A Cooperative Spectrum Sensing Network with Signal Classification Capabilities

A Major Qualifying Project
submitted to the Faculty of
Worcester Polytechnic Institute
in partial fulfillment of the requirements for the
degree of Bachelor of Science

by
Michael Calabro

January 14, 2010

Project Advisor: Professor A. Wyglinski

Abstract

This report describes the design and implementation of the spectrum sensing and signal classification sub-systems of a cooperative network. A sensor blindly receives and calculates the cyclic statistics of a signal decides whether or not the signal represents information or noise. If the statistics of the signal indicate the presence of data, the system attempts to classify its modulation scheme. Finally, the decisions of several independent sensors are combined to provide a reliable estimate of the contents of the spectrum of interest. Independently, sensors correctly classify a signal about 60-70% of the time in a low SNR environment. The data fusion module improves this number significantly - especially as the number of sensors increases.

Acknowledgments

I would like to thank The MathWorks for their financial support on this project. Their support helped purchase several SDR platforms that made verification in hardware possible.

I am indebted to Professor Wyglinski for the opportunities he has given me. Looking back on my experience at WPI, I think I would be hard pressed to find the opportunities that he has given me at any other institution.

To my two teammates Devin & Ishrak: It was a pleasure working with you. Our meetings were fun, and our weekends were interesting.

Executive Summary

Interoperability of wireless equipment is a major concern in modern communications systems - especially in first responder situations where communications infrastructure plays a vital role in the coordination of relief efforts. In New York City on September 11, 2001, and in New Orleans during Hurricane Katrina in 2005, communications infrastructure suffered a complete failure resulting in unnecessary loss of life and diminished rescue capabilities.

Consequently, this has caused the FCC to push for better first responder communications infrastructure. Their primary effort is the P25 public safety communications standard. The FCC mandates that in order to purchase new communications equipment with Homeland Security grants, the equipment must be P25 compliant. This initiative by the FCC to standardize equipment throughout different jurisdictions is an excellent first step toward assuring communications during first responder events.

In addition to legislative and regulatory efforts, technology also offers a solution. As the computational power of wireless devices has increased, modern radio capabilities have expanded dramatically, supporting several simultaneous users, robustness to harsh RF environments, and increasingly high data rates. The next generation of radio equipment will be defined using software code on either programmable logic or microprocessor systems, resulting in software-defined radio (SDR) platforms. These types of wireless devices offer technical solutions to many problems that presented during 9/11 and Katrina.

On 9/11, when the first tower fell, the resulting chaos on the communications channels resulted in a near total breakdown of communications infrastructure. The resulting chaos severely diminished efforts to get responders out of the second tower before it fell and demonstrates the need for a communications system that can organize itself to avoid such chaos [1]. SDRs offer one solution to this problem by dynamically configuring their transmit and receive capabilities. Such capability allows radios to "get out of the way of each other" and could go a long way toward assuring communications during events like this.

SDRs can prioritize data, maximize spectrum utilization, reconfigure their transmit and receive parameters and intelligently form networks among other capabilities [2]. A network defined by these radios is agile, reconfigurable, and robust to hostile operating environments [3].

This technology is still emerging, and it will be some time before it is widely deployed to first responders. In the meantime, the potential capabilities of SDR and cognitive radio present many novel and interesting technical challenges. The Software Defined Radio Forum,

a group made up of industry SDR developers created a challenge to address some of these issues.

The 2009 Smart Radio Challenge, sponsored by the SDR Forum, challenges participants to create a network capable of observing and coordinating first responder radio equipment during a major disaster event [4]. This network represents a proof of concept of the capabilities described above.

This report contributes directly to the WPI team competing in the radio challenge, and specifically addresses the following issues:

- *Spectrum Sensing*: Detecting if a user is operating on a channel. This capabilities avoids unnecessary interference and allows available spectrum to be efficiently utilized [3].
- *Signal Classification*: If a signal is deemed present, this capabilities determines how to talk to it. This allows responders possibly unaware of each others presence to coordinate their activities and maximize their benefit [3].
- *Data Fusion*: Rather than relying on one set of sensor to make decisions, the network relies on many to form high confidence decisions that responders can depend upon [5].

Though such ideas have been around for close to two decades, and were described in detail as early as 1991 [6], only recently has the computational power of current wireless systems been capable of executing the complex mathematics involved in realizing these systems.

This report describes and demonstrates the success in implementing solutions to these three challenges. Signals were blindly classified in white noise with success rates approaching 70% with no prior knowledge of their existence. The data fusion module improved these classification rates to close to 100% as more sensors contributed to a final decision. The classification metrics are comparable to the theoretical and simulated results of other cyclic detectors [7] [8].

These modules were built and demonstrated with MATLAB and Simulink. Where possible, verification was achieved through hardware SDR platforms.

The success demonstrated by these sub-systems establishes the groundwork for the development of an agile network that maximizes spectrum utilization, and is robust to high traffic loads. Additionally, these sub-systems were developed with a black-box approach, such that they are readily deployable with minimal adjustment to other networks that target first responders.

Contents

List of Figures

Chapter 1

Introduction

1.1 Project Motivation

Communication systems interoperability is a major priority in public safety organizations such as firefighters, police services, national guard, paramedics, and other first-responders. This term refers to the ability of different agencies to communicate with each other on their own equipment. Though current systems are often adequate for local first responder events, the infrastructure begins to fail in large-scale events and disaster situations, such as the terrorist attacks on 9/11 and Katrina, when many first responder agencies need to coordinate their efforts in real time. Efforts are underway by federal and state agencies to modernize the public safety communications infrastructure.

Current first responder communication systems rely on legacy analog communications systems. They are very sensitive to network overload and environmental interference [1]. Additionally, equipment made by different manufactures may be incompatible. 9/11 and Katrina highlighted these short comings in the United States. Consequently, the FCC has been encouraging new regulations to standardize first responder communications capabilities [11].

New public safety standards such as Project 25 strive to standardize new communications equipment deployed to first-responder agencies with modern technology that overcomes the limitations of analog equipment. This standardization ensures that agencies will all be using communications equipment capable of talking to each other. P25 compliant equipment must offer an improvement of at least two times the data rate of current analog equipment. Ultimately, efficiency improvements as much as four times are being sought. P25 also seeks to address issues such as prioritization of packets, network roaming, and large, expandable, networks.

In addition to P25, there are efforts within the communications community to raise awareness and demonstrate the potential of modern radio systems. The 2009 Smart Radio Challenge (SRC 2009) sponsored by the SDR Forum represents one such effort [4]. The SRC 2009 seeks to demonstrate the potential of software defined and cognitive radio systems.

1.2 Problem Statement

Catering to the FCC mandate of increased interoperability among emergency responder teams, the SRC 2009 defined the following problem:

> An earthquake has occurred in a major metropolitan area measuring 10.0 on the Richter scale. Existing communications infrastructure is out, and as emergency medical services, police, fire, state, and federal management personnel arrive on the scene from all over the world, they all begin setting up their own communications systems to aid in rescue efforts. As more and more personnel arrive, finding available spectrum becomes a challenge resulting in unintentional interference between communications of various services.

If such a situation occurred in the present, communications infrastructure would suffer a breakdown similar to what happened during 9/11 and Hurricane Katrina in the United States in 2001 and 2005 respectively. To help coordinate first responder teams, the SRC 2009 proposes the creation of an observational network that would facilitate the organization and coordination of radio devices on site. The technical requirement of the challenge was defined as follows:

> Develop a cooperative sensing system that will create and maintain a database of public safety emitters on the scene, including the emitter location, physical layer parameters such as modulation type and transmit frequency, and an association to which emergency team is using this frequency and waveform. There will be at least 20 different emergency response teams present which will be trying to coordinate their activities.

The proposed solution involves many technical challenges. This report solves three of them. The author of this report collaborated closely with two other undergraduates, Devin Kelly and Ishrak Khair, at WPI to deliver a complete network that satisfies the SRC's technical requirement.

In addition to solving each problem theoretically, this network must be complete and at least functioning in software. This means defining and implementing interfaces between different sub-systems. The sub-systems immediately recognized as necessary include a spectrum sensing and signal classification sub-system, data fusion module, geolocation method, a network multiple access scheme, and an environment design.

While considering each sub-system, it was discovered that each problem has many solutions. A major part of this project was considering alternative solutions and making an informed design decision on which solution to pursue. Factors here included time of implementation, simplicity, and effectiveness of the solution.

1.3 Project Definition

This report describes the design and implementation of a cooperative spectrum sensing and signal classification sub-system. Spectrum sensing, signal classification, and data fusion are all important topics in modern wireless communications and are being incorporated as key components in decentralized networks that seek to maximize their computational power and take full advantage of the wireless spectrum.

The SRC 2009 mandates the creation of a cooperative sensing system that will create and maintain a database of public safety emitters during a first responder event. The database should track the geographical location of the emitter, physical layer parameters, and assign a unique identifier to each emitter.

This project demonstrates the implementation of a sub-system capable of determining if modulated information is present on an RF channel. If a signal is detected, it then classifies the signal into one of several P25 mandated transmission schemes. Finally, to ensure a reliable decision is made, the independent decisions of several separate sensors are combined into one decision.

Since this sub-system is part of a larger project, functionality was built with a "black box" approach, i.e. You feed in a set of defined parameters and observe one of several known outcomes.

As a proof of concept, this sub-system will be integrated into one such network, demonstrating the capabilities and potential of this technology. This a larger network whose purpose is to observe a third party communications network, geographically locate the radios on it, and aggregate the data into a persistent database that tracks the users of the third party network. This particular network's objective will be to organize and coordinate first responder teams around a disaster zone, however the applications of this technology are vast.

1.3.1 Spectrum Sensing

Spectrum sensing refers to the process of observing an RF bandwidth to determine if a signal is present on it. Colloquially, this means looking at some spectrum and determining if a signal of interest is present. This is the first step in many applications including dynamic spectrum access and geolocation. If a signal is determined present, the next step is determining how to demodulate its data, this is called signal classification.

1.3.2 Signal Classification

Once the decision as to whether a signal is present on the channel is made, the next step is to classify that signal by its modulation scheme. The approach described in this report relies upon the inherent periodicities introduced to a signal when it is modulated. These periodic, or cyclic, statistics produce observable phenomena on a bi-frequency plane. These cyclic features are specific to different types of modulation, making classification possible.

1.3.3 Data Fusion

Independent sensors are susceptible to making bad decisions. Harsh RF environments can cause sensors to report false positives or false negatives. To increase the reliability of the overall system, a single decision is formed by weighting multiple, independent, decisions made individually by sensors. This process is called data fusion and helps the network make decisions with a high degree of confidence.

1.4 Report Organization

This report is organized into five chapters. The first chapter introduces the project definition and motivation. Chapter 2 explores the different approaches to solving the challenges outlined in Chapter 1. It also provides an introduction to several technical and mathematical topics the reader should be familiar with to understand the implementation details. Chapter

3 lays out the design approach taken by the WPI Smart Radio Challenge team to build the network outlined in the problem statement. This chapter contains lessons learned, descriptions of the design cycles, and the approach the author took to solve his three specific problems. Chapter 4 describes the technical implementation of the three sub-systems this report describes. It provides numerical results and illustrative examples of the topics discussed in chapter 2. The final chapter summarizes the accomplishments of this project and notes several areas where future work is appropriate.

Chapter 2

Fundamentals of Signal Detection, Classification, and Data Fusion

2.1 Digital Communications Primer

Digital communications is a broad field, but this section provides some background to the specific modulation types and terminology that appear frequently in this report. More detailed analysis and communications fundamentals can be found in [12].

2.1.1 Data Modulation

Modulation refers to the way information bits of data are organized. By organizing data well, the entropy of the transmission can be reduced and the information easily recovered by the receiver. This project leveraged the Project 25 (P25) public safety standard to limit the scope of its requirements. Two modulation schemes specified in P25 are quadrature phase shift keying (QPSK) and quaternary frequency shift keying (4-FSK).

QPSK had four constellation points. Each point on the constellation is represented by:

$$S_i(t) = I_i(t)cos(2\pi f_c t) + Q_i sin(2\pi f_c t), \tag{2.1}$$

where I_i and Q_i are the amplitudes of the in-phase and quadrature components of the signal respectively. Figure 2.1.1 shows the structure of a basic QPSK modulator. To decode a QPSK transmission, the incoming signal is multiplied by a cosine and sine. Since cosine and sine are orthogonal to each other, multiplying any in-phase bits with quadrature bits

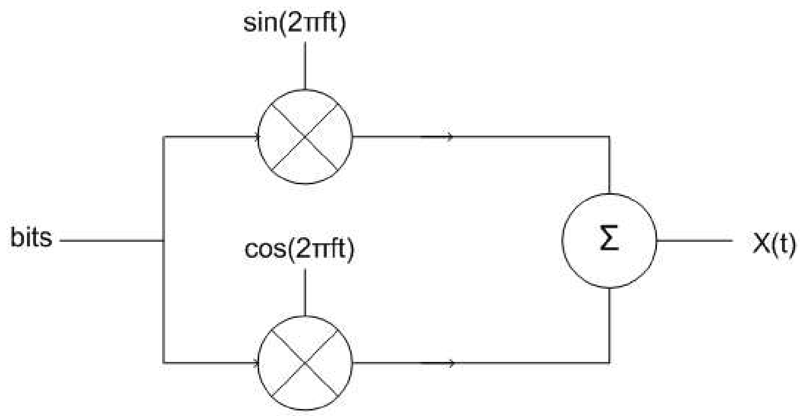

Figure 2.1: Diagram of a QPSK Modulator

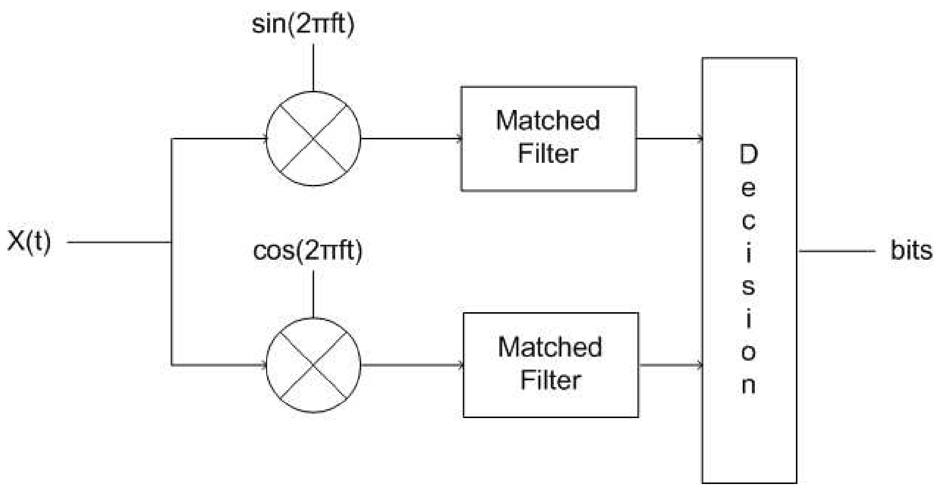

Figure 2.2: QPSK Demodulator with a Matched Filter Receiver Structure

causes the result to be zero. The result of this multiplication is then compared against the constellation and the constellation point that most closely matches the multiplication is chosen.

FSK is a form of frequency modulation that's constellation is represented by discrete frequencies:

$$S_i(t) = M(t)cos(2\pi f_i t), \tag{2.2}$$

where f_i is $f_c + \Delta f$. To modulate and decode M-FSK data, the mixers in Figure 2.1.1 and 2.1.1 are replaced by a bank of M sinusoidal mixers, each oscillating at a distinct frequency.

By definition, modulated data contains inherent periodicities. An illustrative example of this is AM transmission. A message signal $M(t)$ is modulated onto a sinusoidal carrier $sin(2\pi f_c t)$: $X(t) = M(t) * sin(2\pi f_c t)$ After modulation, the output of the modulation has

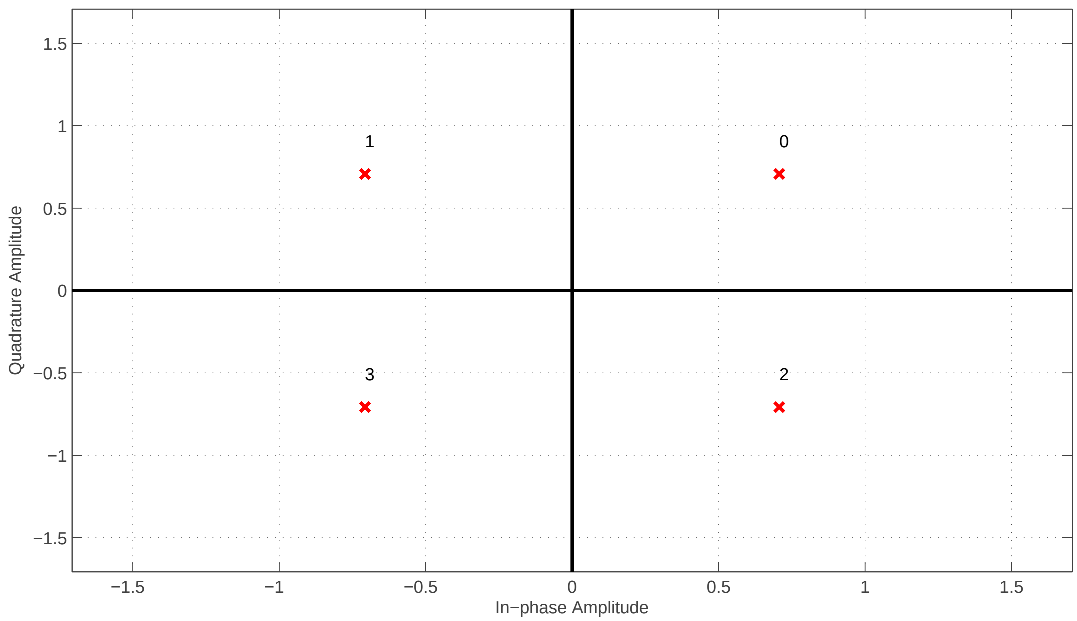

Figure 2.3: Rectangular QPSK Constellation

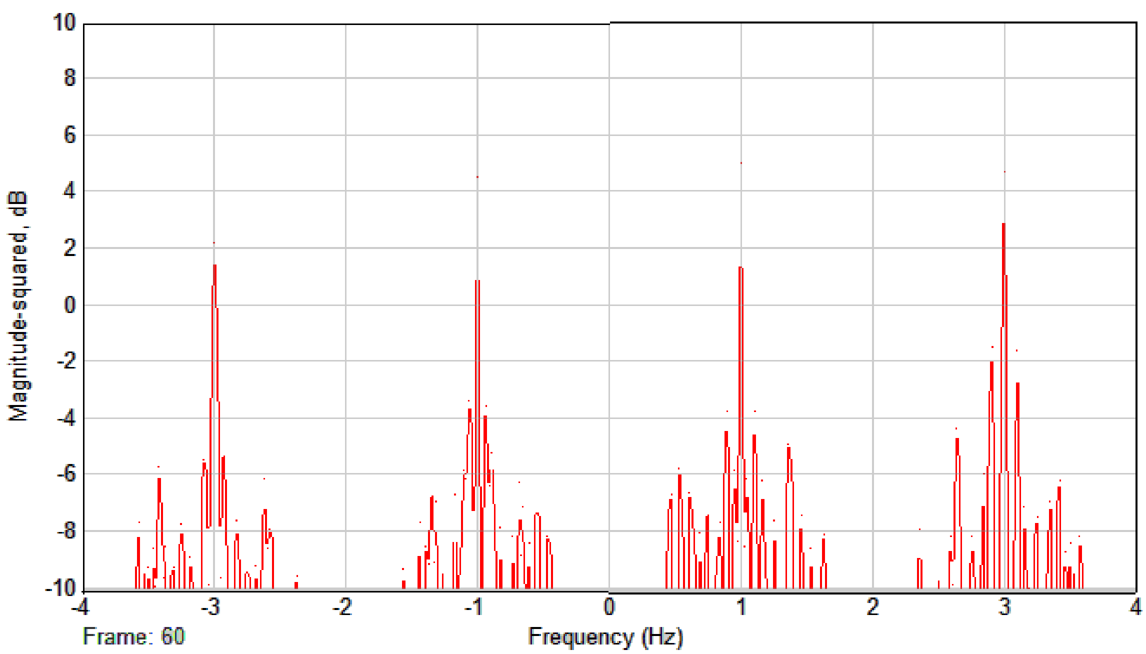

Figure 2.4: 4-FSK Frequency Domain Representation

properties of both the pulse shape (the sinusoid) and the original message.

While these periodicities may not be immediately evident on a spectrum analyzer or oscilloscope, blind (no knowledge of the signal) examination of the cyclic statistics of the signal may reveal statistical features unique to different kinds of modulation. Though this report only deals on extrapolating a modulation scheme from these features, additional parameters such as symbol timing and precise carrier frequency can also be extrapolated.

2.2 Interoperability of Communications Systems

Digital communication systems transmit data through the physical (PHY) layer of a network. The PHY layer includes the modulators, mixers, and other RF equipment that translates voice and data into a format suitable for over the air (OTA) transmission.

The parameters chosen for this hardware include transmits frequency, sampling rates, modulation scheme, data encoding, etc. In order to translate received bits into data, the receiver must be parameterized the same way as the transmitter. The concept of parameterization is well known in software and is becoming an important part of modern radio hardware that implement technical specifications through programmable logic [2].

The term interoperability refers to the ability of two radios to communicate with each other. Much current first responder radio equipment uses analog transmission schemes specific to the manufacturer of the radio. Modern and next-generation first responder equipment is migrating to digital technology with dramatically expanded capabilities and potential [2] [1].

2.2.1 Project 25

Project 25 (P25) is an effort by the Association of Public Safety Communications Officials International (APCO) to standardize modern first responder communications equipment. P25 is still under development and is scheduled for a multi-phase deployment.

Phase 1 is well-defined and is currently being implemented. Phase 1 radios can communicate with legacy analog equipment and modern digital equipment at a channel rate of 9600 bps. This phase standardizes the system infrastructure ensuring that any P25 compliant radio can communicate on the network.

Phase 2 is under development and seeks to improve spectrum utilization. This in-

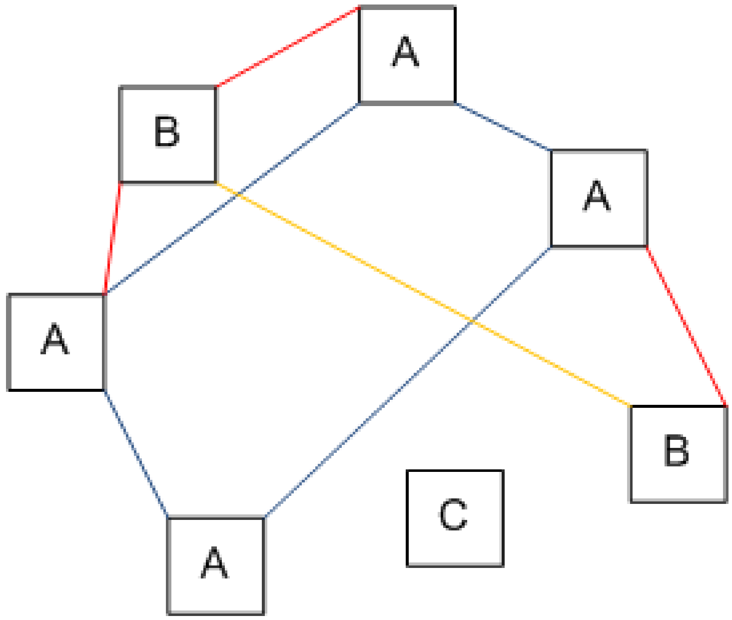

Figure 2.5: The following radios employ communication schemes A, B, and C. A and B are interoperable, C is not. This diagram illustrates the consequences of interoperability on a network and how it influences communications links. In a first responder event, C would be isolated and unable to communicate with other responder teams.

cludes implementing robust roaming and multiple access schemes capable of handling a large amount of users [11].

Once widely deployed P25 infrastructure will support:

- Interoperability within the communications network.
- Digital voice and data transmissions.
- Scalable network scale capable of supporting small or large networks.
- Modern spectrum utilization techniques.

2.3 Software Defined Radio

Figure 2.6 shows a flow diagram of an SDR. Traditionally, the line separating hardware and software was closer to the data sink/source. Modulators, encoders, equalizers, etc., were implemented in hardware. However as high performance programmable logic comes down in cost, radios are increasingly being defined in terms of software.

The sub-systems implemented in this project take advantage of the dynamic nature of SDR platforms. The reconfigurable nature of SDRs make it possible to search out radios operating in a certain frequency band and then determine how to talk to them. The potential for these capabilities extend beyond first responder applications into a variety of consumer applications.

An early application of SDR technology was the SPEAKeasy network developed by the United States military. It allowed digital radios to communicate over a large range of frequencies using various modulation schemes, data encoding methods, and encryption techniques.

A primary goal of implementing SPEAKeasy was to achieve a robust, dynamic communications network that would encourage interoperable communications equipment and provide the information assurance required by the military [2]. Though it was implemented in the 1990's, SPEAKeasy did achieve this goal and demonstrated the promise of future SDR platforms by sending a variety of waveforms through common hardware.

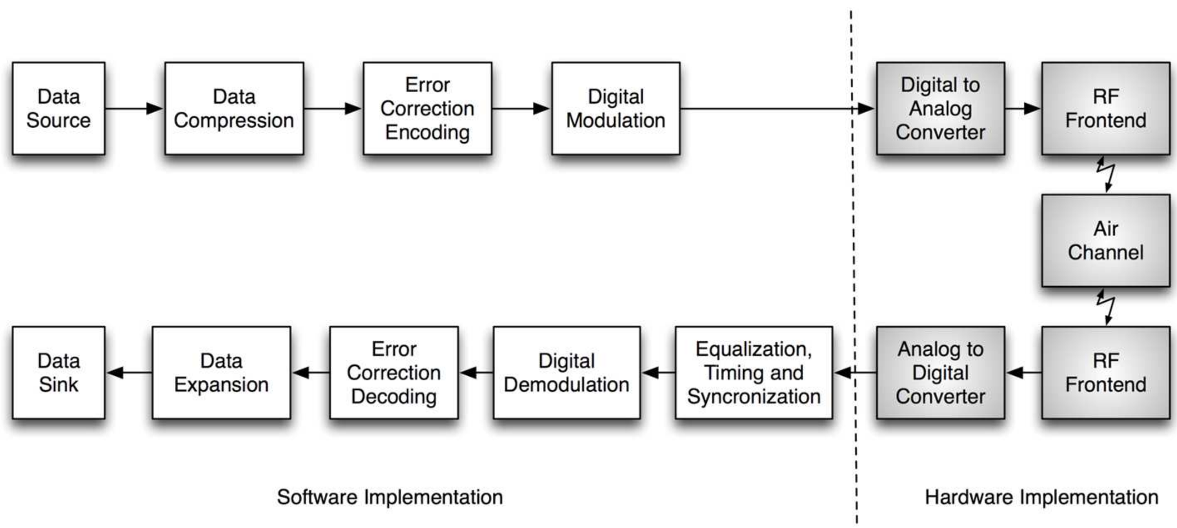

Figure 2.6: Flow diagram showing the evolution of radio implementation from hardware to software.

2.3.1 VITA 49 Radio Transport Standard

The VITA Radio Transport (VRT) standard is a new transport layer protocol that aims to ensure interoperability among SDR platforms. It defines the way digital data and sensor settings are to be transmitted. Embedded within the standard are also ways to configure the SDR radio itself.

An advantage to VRT is that it defines two kinds of packets: data packets and context packets. The data packets of VRT are similar to any other transport layer data packet. What makes VRT unique are the context packets. Most of the fields within these packets are optional, but they offer the potential to transmit all kinds of useful parameters to the radio including, but not limited to: time of transmission, symbol timing, center frequency, sampling times, etc. These packets are capable of defining the physical layer parameters of the radios [9].

Though not directly utilized in this report, VRT packets played an integral role in the geolocation capabilities of this network. Using the context packets, the time of transmission was compared against the time of arrival to form a radial estimate of the transmitters location. Three such estimates were enough to triangulate a reasonable estimate of the radios actual location.

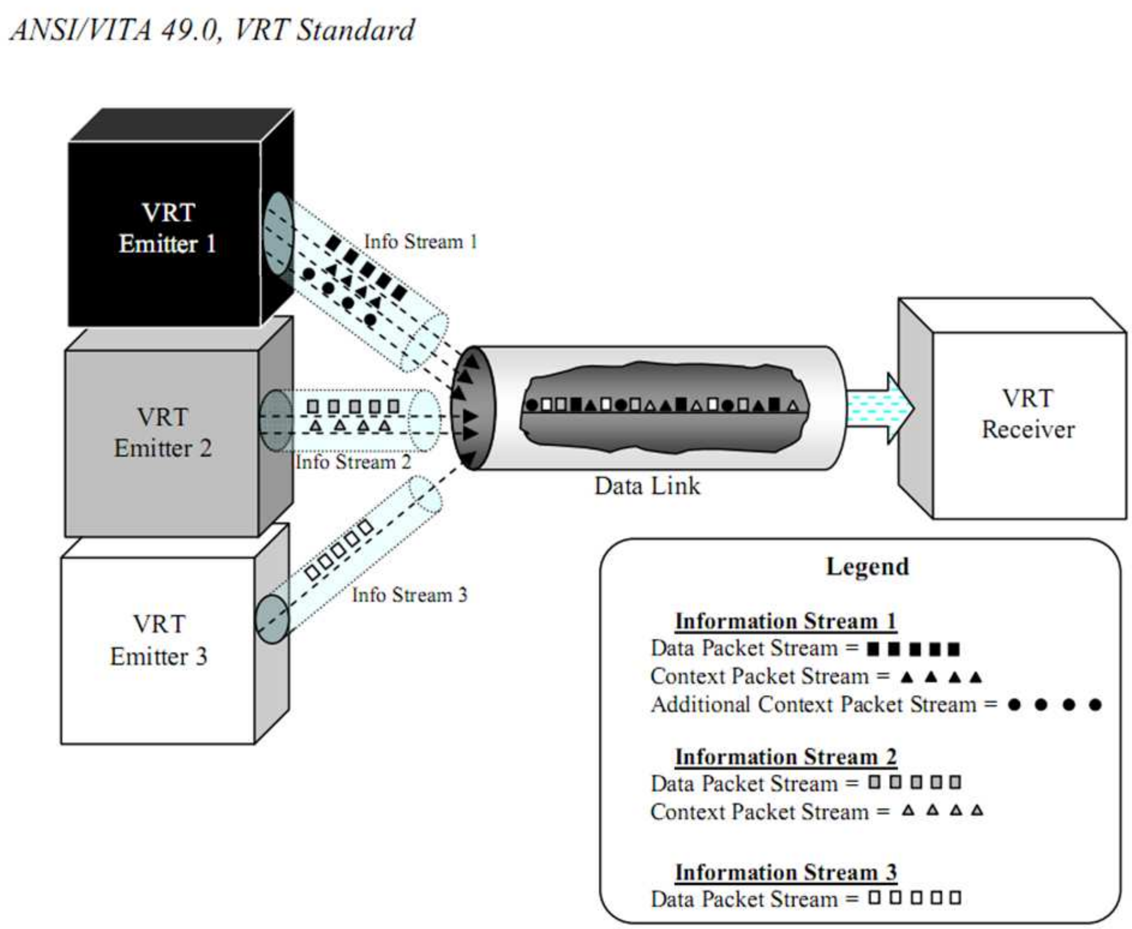

Figure 2.7: This diagram shows the packet and flow structure of a VRT communications link. Contextual packets are interleaved with data packets. The contextual packets are transmitted at a slower rate than the data packets. [9]

2.3.2 The Universal Software Radio Peripheral 2.0 (USRP 2.0)

The USRP 2.0 is an SDR platform developed by Ettus Research. Compared to other SDR platforms it is relatively cheap (about $1,400, up to $2,000 with daughterboards and antennas as of late 2009).

The USRP 2.0 is modular in nature. At its core is a motherboard which contains 4 ADCs and 4 DACs with digital I/O lines that attach to user-selected daughter cards. These daughter cards act as RF frontends for the radio and determine the frequency range capability of the radio. A major advantage of this platform is its ability to interface directly with GNU radio.

GNU radio is an open source software project that encourages the development of software libraries and packages for the USRP. This interface with software allows the RF frontend to be tuned through software. Configurable parameters include gain, frequency, interpolation rates, and decimation rates.

This report utilizes an experimental interface between the USRP platform and Simulink developed by the WPI Wireless Innovation laboratory. This interface allows a model in Simulink to control the physical radio. Where it was possible and feasible, some of the algorithms implemented in this report were verified and tested on data transmitted through this SDR platform using this Simulink interface.

2.3.3 Cognitive Radio

A cognitive radio is defined in [13] as:

> An intelligent wireless communication system that is aware of its surrounding environment, and uses the methodology of understanding-by-building to learn from the environment and adapt its internal states to statistical variations in the incoming RF stimuli by making corresponding changes in certain operating parameters (e.g. transmit-power, carrier-frequency, and modulation strategy) in real-time, with the two primary objectives of highly reliable communications whenever and wherever needed, and efficient utilization of the radio spectrum.

The fundamental term inferred in this description is adaptability. Cognitive radios are a natural extension of software defined radios (SDR), or radios that have components classically implemented in hardware implemented in programmable logic.

As a layer on top of SDR, the term cognitive radio is inclusive of the capabilities defined in the SPEAKeasy project [2]. A radio that can re-define its physical layer parameters becomes quite versatile. Intuitive applications of this technology include dynamic spectrum access to maximize spectrum utilization, cross-standard radios, and signals intelligence.

The 802.22 Wireless Regional Area Network (WRAN) standard is currently under development and utilizes cognitive radio technologies extensively. One goal of 802.22 is to use dynamic spectrum access to alleviate spectral congestion in the TV bands in densely populated areas. The spectrum sensing and signal classification capabilities of this network are integral parts of these applications [14].

2.4 Fundamentals of Signal Detection

Two reasons to sense spectrum are: Primary signal detection, and spectrum opportunity detection (SOD). SOD refers to the detection of a spectrum opportunity in dynamic spectrum access and is outside the scope of this paper. We will be focusing on detecting primary signals and then classifying their physical layer parameters.

2.4.1 Energy Detector

Energy detectors are based on hypothesis testing. A hypothesis test validates or rejects an assumption by comparing a test statistic to a threshold that represents the null hypothesis. A basic hypothesis model of a signal in an AWGN channel is given by [12]:

$$H_0 : y(t) = n(t), \tag{2.3}$$

where no signal is present and n(t) is the channel noise, and:

$$H_1 : y(t) = x(t) + n(t), \tag{2.4}$$

where the signal is present in additive noise.

This approach relies only on the energy present in the channel. Since the energy of a signal is defined as $\int^{\{}\infty\}_{\{}\infty\}\{|f(t)|^2, dt\}$, no phase information is required. The underlying assumption is that with the presence of a signal in the channel, there would be significantly more energy than if there was no signal present.

Within this premise lie two major disadvantages of energy detectors. First is that by definition, no signal will be detected that is below the noise floor. Second, the detector is susceptible to two types of errors.

A type I error is when a signal is reported detected, but in fact it was actually a deceptively noisy channel causing a false positive. A type II error is made when the signal is buried in the noise and is mistaken for noise. [15] shows that missed detections can occur as often as 40% and that false positives as often as 30%. Both of these estimates are conservative and unacceptable for a network of this nature.

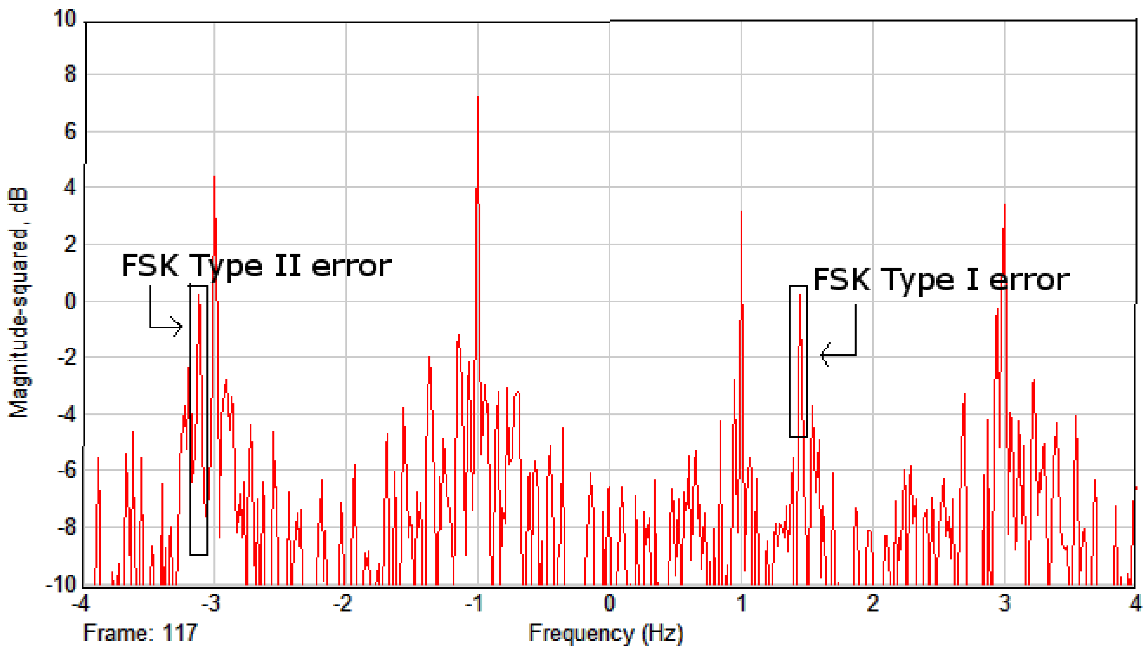

Figure 2.8: Possible Type I and Type II errors of 4-FSK detection using an energy detector.

2.4.2 Cyclostationary Analysis

This report relies upon the second order cyclic statistics of a signal to sense and classify a signal. Higher order cyclic statistics are also quite useful, and are necessary to differentiate similar modulation schemes (4-QAM versus 16 QAM for example), but are outside the scope of this paper.

A cyclostationary signal is a signal whose statistics vary periodically with time. By definition a signal x(t) is wide-sense (meaning only the first and second moments do not vary with respect to time) cyclostationary if its mean and autocorrelation are periodic [16]:

$$R_x(t,\tau) = R_x(t+T_0,\tau)\forall t,\tau, and M_x(t+T_0) = M_x(t). \tag{2.5}$$

The periodic nature of this signal allows it to expressed as the Fourier series:

$$R_x(t - \frac{\tau}{2}, t + \frac{\tau}{2}) = \sum_{\alpha} R_x^{\alpha} \tau e^{-j2\pi\alpha t}, \tag{2.6}$$

where the cyclic frequency $\alpha = m/T_0$. The Fourier series is the decomposition of a signal into a summation of contributing frequencies. Expanding Equation (2.6) yields the cyclic autocorrelation function (CAC):

$$R_x^{\alpha}(\tau) = \lim_{T\to\infty} \frac{1}{T} \int_T x(t + \frac{\tau}{2}) x(t - \frac{\tau}{2})^* e^{-j2\pi\alpha t}\, dt \tag{2.7}$$

It is important to note that at $\alpha = 0$ Equation (2.7) is equal to the traditional autocorrelation of the signal. This can act as an important sanity check when verifying an implementation. The CAC is the first step towards obtaining the cyclic statistics of the signal. The spectral correlation function (SCF) can be obtained by the Fourier transform of $x(t)$ [6] and is given by:

$$S_x^{\alpha}(t, f) = \int_{-\infty}^{\infty} R_x \alpha(\tau) e^{-j2\pi f \tau}\, d\tau \tag{2.8}$$

The SCF is a cross correlation between frequency components of the signal separated by $f + \frac{\{}{\alpha}\}\{2\} and f - \frac{\{}{\alpha}\}\{2\}$. Since practice dictates that only a finite window of samples will be available for processing, the SCF must be estimated from these samples. This introduces computational trade offs to the algorithm that will be discussed later. The cyclic periodogram is defined as:

$$S_{x_T}^{\alpha}(t, f) = \frac{1}{T} X_T(t, f + \frac{\alpha}{2}) X_T^*(t, f - \frac{\alpha}{2}), \tag{2.9}$$

where X_T is the time varying Fourier transform defined in [16] as:

$$X_T(t, f) = \int_{t-\frac{T}{2}}^{t+\frac{T}{2}} x(u) e^{-j2\pi f u}\, du. \tag{2.10}$$

Finally, we are interested in computing the correlation coefficients of the SCF (referred to in [16] as the spectral coherence function (SOF)). The magnitude of the SOF varies from 0 to 1 and represents strength of second order periodicity within the signal. The SOF is given by:

$$C_{x_T}^{\alpha}(f) = \frac{S_x^{\alpha}(f)}{\sqrt{S(f + \frac{\alpha}{2}) S(f - \frac{\alpha}{2})}} \tag{2.11}$$

The SOF contains the spectral features of interest. These features are non-zero frequency components of the signal at various cyclic frequencies. Different modulation schemes contain spectral components at different cyclic frequencies. For example QPSK may contain a cyclic frequency component at $\alpha = 0.1 * Fs$ while BPSK contains cyclic components at $\alpha = 0.4 * Fs and \alpha = 0.6 * Fs$. These distinctions allow signals to be classified from cyclic analysis.

2.4.3 Signal Classification as an Extension of Detection

Approaches to signal classification can be classified as likelihood based or feature based. Likelihood based classifiers assume prior knowledge or a good estimate of the signal and channel statistics. When trying to accomplish blind detection, this is not feasible, and the consequence is a poorly performing detector.

Unlike likelihood based tests, feature tests look for key statistics in the received signal and form a decision based on these estimates. Traditional feature based approaches also relied heavily upon prior knowledge of the signal. Though they are technically sub-optimal, because they are only looking for specific features their computational complexity is considerably lower than a likelihood based approach [7]. However, their reliance on prior knowledge of the signal (symbol timing, carrier phase, timing offset, etc.) make them ineffective.

Cyclostationary detectors have a major advantage over any other detection scheme: they draw their conclusions from the periodic statistics of a signal. These statistics are calculated at the receiver and assume no prior knowledge of the transmission. The advantage here is that with very little additional processing, a signal can be classified with a high degree of reliability [6].

When examined on the bi-frequency plane created by Equations (2.8) and (2.11), each kind of digital modulation has distinct features. These features can be extracted to create a profile as proposed in [16]. These profiles are simple to create and require no additional mathematical derivations - they are simply snapshots of Equation (2.11).

The two modulation schemes of interest to this report are 4-FSK and QPSK. They are distinct enough not to require any additional processing on their cyclic profiles to differentiate them. To differentiate between similar modulation schemes like 4-QAM, 16-QAM, and 64-QAM, additional computation is needed. This is because similar modulation schemes exhibit cyclic frequency components on the same cyclic frequencies, and their second order statistics look quite similar. With additional processing, they can be differentiated by computing the cyclic cumulants of the signal (related to the statistics moments) [7].

Another important topic to mention is that of excess bandwidth. In digital communications when a signal is transmitted, it is passed through a pulse-shaping filter to reduce intersymbol interference. The rolloff factor, β of this filter introduces redundancy into the

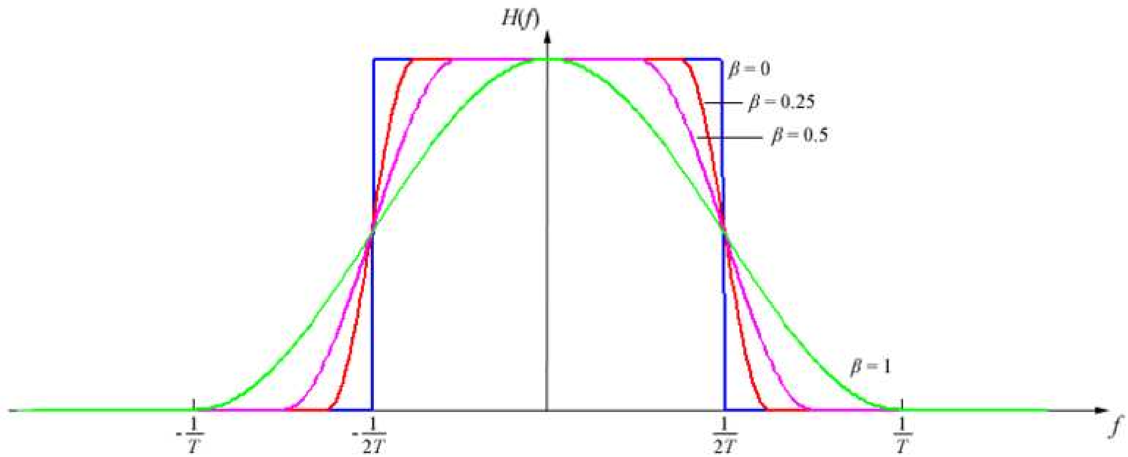

Figure 2.9: Effect of different β values on the frequency response of a pulse shaping filter. [10]

signal. Figure 2.4.3 illustrates the consequences of different β values. The purpose of this redundancy is to act as a guard against ISI and channel distortion. The more excess bandwidth introduced by the pulse-shaping filter, the more redundancy is in the signal, consequently, the clearer the spectral features are.

If $\beta = 0$ there is no excess bandwidth and no redundancy in the transmission, so spectral features are impossible to distinguish from noise. Realistic β values of 0.3 and 0.4 allow for the detection of spectral features. In harsh channel conditions, manipulating this value might achieve better classification results.

The capability to accomplish this signal classification as a natural extension of the cyclic detector is a major advantage of cyclostationary signal analysis. Cyclic detectors are robust to multipath, white noise, and can practically function with little to no prior information about the signal of interest [7].

2.5 Data Fusion Approaches

Once independent decisions have been made on the nature of the signal, combining the decisions into one reliable estimate can be done through several classical approaches to data fusion. The three approaches examined in this report are: voting, weighted averaging, and Bayesian filtering.

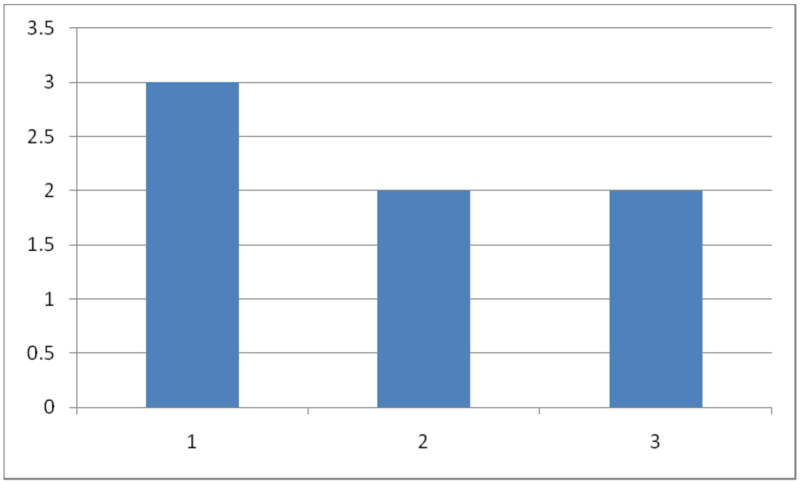

Figure 2.10: An example of a ballot box voting system where a simple majority decides the outcome.

2.5.1 A Voting System

A voting system is the simplest fusion system to implement. Out of the three candidate decisions, the data fusion block counts the number of sensors that "vote" or independently decide on each solution. A majority vote would decide on the output of the data fusion block.

For example, if three sensors input 1, two input 2, and two input 3, the output of the fusion system would be 1. Though a non-optimal solution, it is extremely simple. Figure 2.5.1 shows this. It could be made more reliable by requiring a two thirds majority, or even a four fifths majority [5].

2.5.2 Weighted Average

Here each sensor would output two numbers, its "decision" as defined above, and an additional number ranging from 0 to 1 which represents a self-assessment of its reliability. If the sensor determines that it is in a hostile environment, it would choose a low value. If the environment is determined to be friendly, it would choose a high value. These determinations could be made by observing the bit error rate of a known sequence transmitted between the base station and sensors to maintain order and synchronization within the network.

An additional factor to consider is the probability of false detection. If it can be empirically shown that this probability can be directly correlated to the channel environment, then depending on the environment, this could also be factored in as an additional weighting factor [7].

Though it may be difficult to determine optimal values for these weights, tunable estimates may be obtained by calculating channel characteristics when no signal is present and when a known signal is present [17] [7].

2.5.3 Bayesian Filtering

The most complex solution, Bayesian filtering could offer the most reliable results. Given the measurement z and the true state x, the probability of a given state, x, is given by:

$$P(x_k \| Z_k) = \frac{P(z_k \| x_k) P(x_k \| Z_{k-1})}{P(z_k \| Z_{k-1}}), \tag{2.12}$$

One type of filter being considered is the Kalman filter, described in [17], which is effective at estimating a linear system from a series of noisy measurements. The Kalman filter is recursive in nature - meaning that it requires the most recent estimate and the current measurement to compute the next estimate.

Given an initial predicted state, the Kalman filter projects the anticipated state $\hat{x}^-$,

$$\hat{x} = A\hat{x}_{k-1} + Bu_{k-1}, \tag{2.13}$$

where A and B represent the relationship between the a priori estimate and the current time step, k, and the error covariance P_k^- given by:

$$P_k^- = AP_{k-1}A^T + Q, \tag{2.14}$$

respectively. After computing these priori estimates, the filter using an actual measurement to weight its estimate.

$$K_k = P_k^- H^T (HP_k^- H^T + R)^{-1} \tag{2.15}$$

Next, the estimate of the state, $\hat{x}$, is revised based on the measured data, z_k,

$$\hat{x} = \hat{x}_k^- + K_k(z_k - H\hat{x}_k^-, \tag{2.16}$$

Finally, the error covariance, P_k is updated.

$$P_k = (I - K_k H)P_k^- \tag{2.17}$$

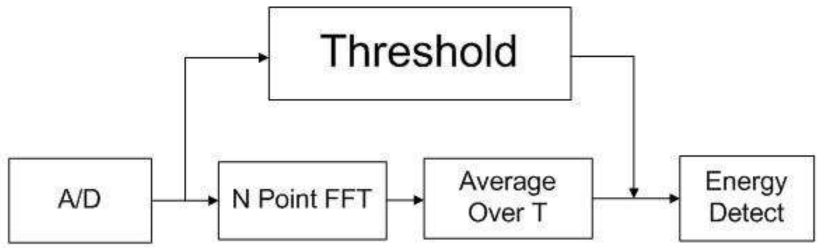

Figure 2.11: Block diagram of an energy detector

The filter is computing a state estimate based on an a priori estimate, $\hat{x}_k^-$ and a difference weighted by K between the actual measurement z_k. In the above equations, R and Q represent the measurement and process noise respectively. R is pre-computed with the known condition that no signal is in the channel. Since the modulation schemes of the expected signals are known, Q can also be estimated. The minus sign in $\hat{x}^-$ denotes that $\hat{x}$ is an a priori estimate of the state.

2.6 Literature Review & State of the Art

2.6.1 Signal Detection & Classification

Energy Detection

Section 2.4.1 describes the technical implementation of an energy detectors. Reference [15] discusses the usefulness of energy detectors as they pertain to surveying spectrum. Energy detectors are capable of providing a fair assessment of spectrum utilization and require only magnitude information of the signal.

One significant drawback of energy detectors is their inability to distinguish between false positives and true signals. Reference [8] examines this issue as it relates to the spectrum surveying conducted in [15].

Energy detectors have significant drawbacks when being used to detect primary signals of interest. They cannot distinguish between desired signals and noise, cannot detect signals below the noise floor, and cannot distinguish between primary and secondary users. References [15] [8] both agree that cyclic detectors offer significantly better performance over an energy detector, but acknowledge the simplicity of energy detectors as an important advantage when considering a spectrum sensing solution.

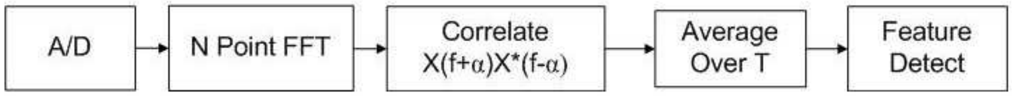

Figure 2.12: Block diagram of a cyclic detector.

Cyclostationary Detection

Cyclostationary feature detection offers better performance than an energy detector and an intuitive extension into signal classification. It requires both phase and magnitude information of the signal.

Since this method relies on the cyclic statistics of the signal, it is computationally intensive. Reference [18] describes two computationally efficient algorithms that provide low resolution estimates of the cyclic spectrum. The Strip Spectral Correlation Algorithm (SSCA) and FFT Accumulation Method (FAM) are both far more efficient than computing the raw cyclic periodogram described in Equation (2.8).

The disadvantage to these algorithms as noted by [7] is that the estimates they provide are often too low a resolution to distinguish spectral features of interest. Reference [7] proposes a method that uses the SSCA to estimate cyclic features of interest and then computes an abbreviated version of the cyclic periodogram around that cyclic frequency to flesh out the spectral features of interest.

Reference [6] offers a detailed overview of the mathematical techniques employed in cyclostationary signal analysis.

A natural extension of cyclostationary feature detection is signal classification. Though [6] touches upon this possibility, [16] and [7] go into detail on the subject. Both [16] and [7] use some form of Bayesian filtering to classify signals. Their signal classifiers are capable of covering a large range of modulation schemes and differentiating between similar ones.

There is no analogy to this signal classification capability in energy detection, making it a key advantage cyclic detectors have over energy detectors.

2.6.2 Data Fusion

A voting system describes a solution where the final output decision is based on choosing the input decision that the most sensors decided was the most probable output. An illustrative analogy is a political ballot question: "Is this a good idea?" If a simple majority is required to pass the measure and five people choose yes and twenty people choose no, then the decision is: "No, this is not a good idea."

Weighted averaging improves upon the voting system by accounting for the reliability of each sensors decision. That is if five highly unreliable sensors report the same decision and one reliable sensor reports a different decision, the decision of the one reliable sensor is chosen as the most reliable decision.

Reference [5] offers a detailed analysis of the challenge of reliable data fusion. Compared to many of the fusion problems described in [5], the fusion problem in this network is relatively straight forward.

2.7 Chapter Summary

This chapter compares and contrasts the various technical approaches to solving the problems of signal detection, signal classification, and data fusion. Additionally, it provides the mathematical and qualitative background necessary to understand the implementation details of the proposed solution described in this paper. Finally, it contains a survey of the literature around these topics and offers up several excellent sources where additional detail into a specific topic can be found.

Chapter 3

Design Approach

This section describes the engineering approach taken to design a solution that would achieve the stated goals of the 2009 Smart Radio Challenge. Though this report describes the technical implementation of specific sub-systems, the overall network architecture will be described for completeness.

For completeness, the design challenge is reiterated:

> Develop a cooperative sensing system that will create and maintain a database of public safety emitters on the scene, including the emitter location, physical layer parameters such as modulation type and transmit frequency, and an association to which emergency team is using this frequency and waveform. There will be at least 20 different emergency response teams present which will be trying to coordinate their activities.

To solve the problem, our team took a divide-and-conquer systems engineering approach followed by systematic integration. We examined the system as a whole and defined various functions as black-box sub-systems that could easily be pieced together.

3.1 Systems Engineering

Our design approach followed an iterative model similar to Figure 3.1. The fundamental advantage of such a process is that it allows the developer to apply the lessons learned on previous iterations to new iterations. Since the challenges facing us were complex in nature, we were able to leverage this advantage as a huge benefit. The trade off was an extended development time, but the advantage in understanding it offered made it worthwhile.

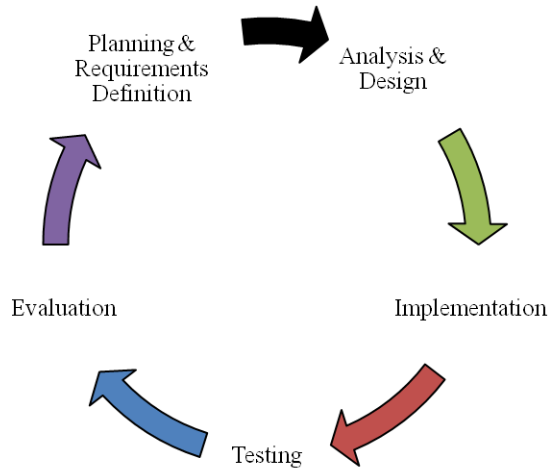

Figure 3.1: The iterative design process.

3.1.1 Initial Planning

The first step in the design of this network was the definition of the specific technical challenges involved in creating a network of this scope. Applying the definition of the SRC 2009 challenge problem, it was determined that the network needs to implement technical solutions to the following problems:

- *Spectrum Sensing*: Scan an RF bandwidth of interest & determine if modulated data is present on the channel. This process will need to be robust to false positives and false negatives.
- *Signal Classification*: If a signal is present, its PHY parameters. Signal classification defines the parameters that other radios need to talk to the radio being classified. This allows for greater organization throughout the whole network.
- *Data Fusion*: Combine data from several sources and form one decision with a high degree of confidence. Independent sensors can make poor decisions if they are operating in a harsh environment. However, a single aggregator, combining the decisions of multiple sensors can provide a highly reliable estimate.
- *Geolocation*: During any first responder event, it is desirable to know the exact location of all personnel. The communication equipment used by first responders offers an intuitive way to locate them. Persistently tracking the location of first responder transmitters allows for geographical guidance to be provided and assistance to be dispatched if necessary.

- *Channel Modeling*: The network must be capable of operating in a variety of first responder scenarios. These include events that occur in rural, suburban, urban, and underground environments. Modeling the channels correctly is essential to assuring the flow of information during crises.
- *Network Synchronization & Multiple Access*: Coordinating the individual sensors of the network is necessary to achieve a cooperative network. An intuitive example are multiple sensors scanning different RF bandwidth, trying to detect a signal. Independently, their estimates are unreliable and difficult to aggregate. By organizing the communication of the network in an effective manner, more reliable estimates are obtained and the network will be more dynamic.

Defining the design requirements was a result of breaking down the problem statement sentence by sentence and determining qualitatively what the sentence was asking for. Once we knew what it wanted (for example the geographical location of a radio), we took this requirement and determined what the technical requirements for it were and what assumptions if any we would need to make.

Fleshing out exactly what each term meant involved a lot of research on the behalf of all team members. The total time to develop and understand the implications of these requirements was about a month.

Once the design requirements were fleshed out in detail, responsibilities for individual team members were established and a schedule was developed to define the development cycle of the network.

Figure 3.2 reflects the schedule defined for the spectrum sensing, signal classification, and data fusion sub-systems. A few weeks are unaccounted for at the end of the schedule to account for unexpected delays. These sub-systems are naturally suited to an iterative design approach. The signal classification sub-system builds directly off of the sensing sub-system, and the data fusion sub-system handles the outputs of the signal classification module.

This schedule was fairly closely followed and the milestones were all accomplished +/- a week of their intended completion date. This was accomplished by being realistic in understanding the challenges from the beginning of the project. Realizing a system in software and/or hardware can take a lot of time. The most important lesson learned in progressing through the implementation stages was to stick to the iterative model. This model ensures you do not get caught up in pedantic theoretical details and forces you to grasp key concepts. As a solution to the problem evolves, most of the theoretical nuances become trivial and easily understood.

ID	Task Name	Start	Finish	Duration
1	Initial Planning	4/10/2009	5/11/2009	22d
2	**Cycle 1**	**4/16/2009**	**5/28/2009**	**31d**
3	**Spectrum Sensing**	**4/16/2009**	**5/28/2009**	**31d**
4	Research Methodologies	4/16/2009	4/29/2009	10d
5	Define Approach Trade-offs	4/27/2009	5/28/2009	24d
6	**Cycle 2**	**5/21/2009**	**8/28/2009**	**72d**
7	**Spectrum Sensing**	**5/21/2009**	**8/28/2009**	**72d**
8	Determine Solution	5/21/2009	6/2/2009	9d
9	Implement Algorithm	5/27/2009	8/7/2009	53d
10	Verify	7/13/2009	8/28/2009	35d
11	**Signal Classification**	**7/23/2009**	**8/28/2009**	**27d**
12	Research Methodologies	7/23/2009	8/28/2009	27d
13	**Final Cycle**	**8/31/2009**	**11/25/2009**	**63d**
14	**Spectrum Sensing**	**8/31/2009**	**10/6/2009**	**27d**
15	Optimize Parameters	8/31/2009	9/21/2009	16d
16	Optimize Code	9/3/2009	9/28/2009	18d
17	Additional Verification	9/14/2009	10/6/2009	17d
18	**Signal Classification**	**9/22/2009**	**10/29/2009**	**28d**
19	Implement Algorithm	9/22/2009	10/16/2009	19d
20	Verify	10/8/2009	10/29/2009	16d
21	**Data Fusion**	**10/19/2009**	**11/25/2009**	**28d**
22	Determine Weight Metrics	10/19/2009	11/3/2009	12d
23	Implement Algorithm	10/29/2009	11/17/2009	14d
24	Verify Performance Increase	11/10/2009	11/25/2009	12d

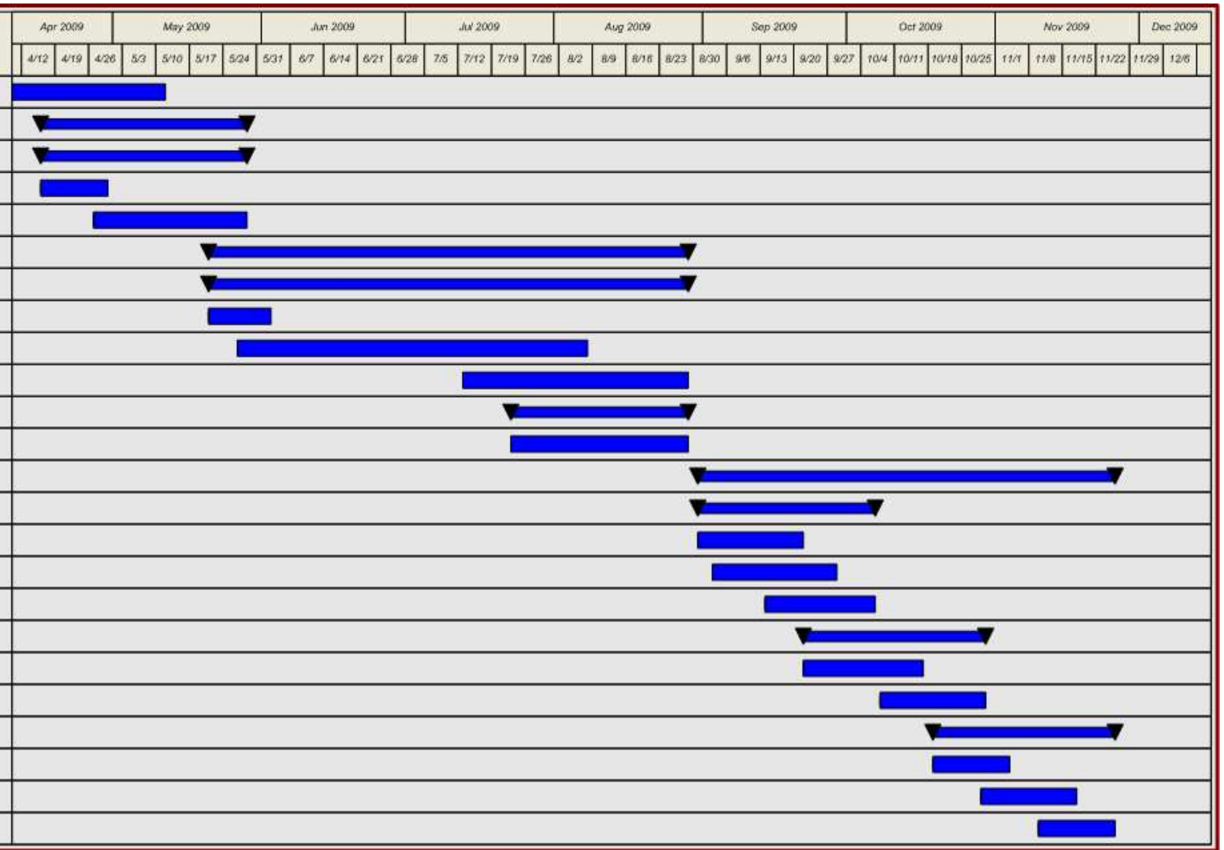

Figure 3.2: Gantt Chart showing the actual work schedule followed

3.1.2 Design Cycle 1

The first design cycle was dominated by research into different spectrum sensing approaches.

Requirements

The primary objective at the end of the first development cycle was to have a clear idea how to proceed with the spectrum sensing sub-system. By extension, this means understanding the trade offs between various approaches and having at least a high level grasp of the different approaches.

Regarding the network as a whole, this cycle focused on defining the functional blocks of the network. Due to the nature of the requirements, a black-box setup was decided upon. The goal of this decision was to minimize the amount of re-design that would need to be done in the event of a major setback.

This decision began paying dividends quite early in the design process (July '09) when the entire problem was redefined. Due to a misunderstanding of the nature of the challenge, the entire network architecture to this point needed to be reconsidered. The team's initial belief was that it was their task to design and implement a next generation communications network infrastructure for first responders. The scope of this problem was far larger than the observational network actually defined.

When we realized our error, it was a relief to know that the network re-design would not mean any additional functionality (in fact it meant far less). Figure 3.3 shows the division of work by team member.

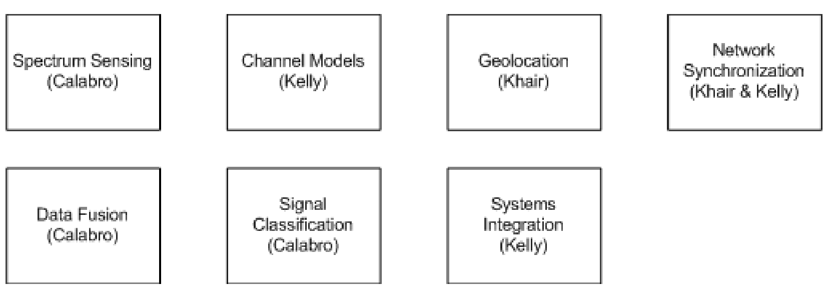

Figure 3.3: This diagram shows the black box nature of the network and the team member assigned to complete each functional block.

Design & Implementation

Research revealed two popular ways to accomplish spectrum sensing. [15] proposes a method based on the energy of the envelope of a signal. This energy detector works by setting a threshold. If energy levels pass this defined threshold, the detector concludes that a signal is present. Below the threshold, the conclusion is that there is no signal present.

While simple to implement and easy to understand, several flaws become immediately apparent on further inspection.

1. Having an energy threshold as the metric for signal detection leaves the system vulnerable to excessive false positives or false negatives from noise.

2. If a signal is determined to be present, no further information is offered about it by the detector other than "it exists."

Reference [7] proposes another solution that relies on the statistics of the input signal to make a decision. This approach calculates the second order cyclic statistics of a signal and based upon the statistical features on a bi-frequency plane of frequency, f and cyclic frequency, α determines if a signal is present. Signals with non-zero components at $\alpha \neq 0$ may be called cyclostationary.

This method relies upon the fact that modulated data is a cyclostationary process and noise is not. Furthermore, the statistical features that are produced in the analysis are specific to the modulation scheme that was used to encode the data. This approach offers two key advantages over energy detection:

1. By Equation (2.5), noise is not a cyclostationary process.

2. The spectral features revealed in the statistical analysis can be used to gather further information about the signal - including its PHY layer parameters.

Other than robustness to noise and additional information, the other primary trade off to consider is complexity. Though its benefits are alluring, cyclostationary signal analysis involves deep math and potentially very high algorithmic complexity. Despite this, a design decision was made to pursue a cycle detector solution because of its natural extension into the signal classification module.

3.1.3 Design Cycle 2

Requirements

The second cycle centered on choosing a spectrum sensing solution and implementing it. Though it is the most complex solution, cyclostationary analysis was chosen. Cyclic analysis was the preferred solution because of its natural extension into signal classification and robustness to noise. Also, due to the simple nature of the P25 public safety standard, it would not need to differentiate between similar modulation schemes, removing a layer of complexity.

Design & Implementation

The first efforts to implement a multi-cycle detector followed [16]. Though the results in [16] looked promising, the first attempt at implementation produced plots similar to figures 3.4 and 3.5.

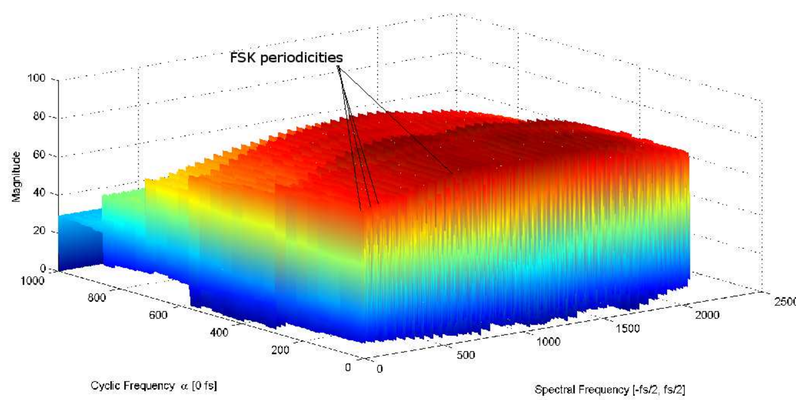

Figure 3.4: SCF of FSK - first attempt

These errors were the result of an incorrect understanding of the cyclic frequency α. In order to understand the concept of a bi-frequency plane, the more rigorous papers by William Gardner were useful. In [6] Gardner derives from first principles a multi-cycle detector. Regarding α, the key to understanding it, for this author, was to consider each cyclic frequency component of the signal as a Fourier series.

A non-zero frequency component at $\alpha! = 0$ indicated some amount of periodicity in the signal.

As the author was coming to grips with the mathematics surrounding periodic statistics,

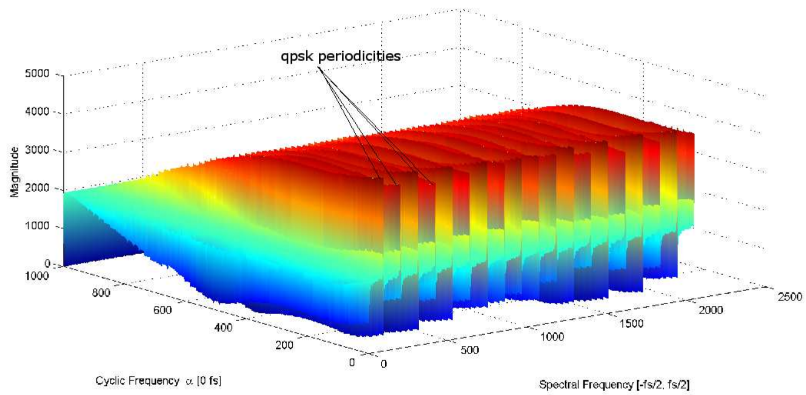

Figure 3.5: SCF of QPSK - first attempt

a lot of work was done by the SRC team as a whole to define the architecture of the network they were building. A hierarchical approach to the network was decided upon for simplicity. That is information would be gathered at a bottom layer of the network and as it moved up through the network, decisions would be made based on that information. Ultimately at the top layer of the network would reside the database containing the responder team data. Figure 3.6 shows the basic functionality of the network.

Assessment

At the end of this cycle the network vision was well defined. The spectrum sensing and signal classification sub-systems were also nearing completion. The biggest challenges during this cycle stemmed from the complex mathematics that define periodic statistics. The biggest take away was that first principles can make some of the most daunting math far more manageable.

3.1.4 Exit Cycle & Systems Integration

Requirements

As this cycle began, the spectrum sensing and signal classification modules were mostly complete. Remaining work to be done on them included optimizing the code for speed and making the code as EML compliant as possible.

Since this was the last design cycle, all of the sub-systems that the author was responsible for needed to be functioning and ready to be integrated into the network. The last remaining technical task was to implement a data fusion scheme.

Design & Implementation

The preferred solution for data fusion in this situation was a weighted average of the data. Each spectrum sensor, in addition to outputting the estimated channel state, would add in a self-calculated reliability factor. The reliability factor is calculated from the perceived noisiness of the channel and the reliability of the communications between the sensor and the base station. As the channel noise and error rate between the base station and sensor increase, the reliability factor decreases.

When implementing the data fusion block, the best performance was achieved by empirically tweaking how the reliability factor was calculated. The reason weighted averaging was chosen over the more accurate Bayesian filtering was the simplicity of the implementation. With the design phase of this network ending and a desire to move on to the other deliverable set forth by the SDR forum, time of implementation was a considerable factor in this decision.

Assessment

As this phase ended, two blocks had been created: a spectrum sensing & signal classification block, and a data fusion block. With well defined inputs and outputs, these two blocks can be seamlessly integrated into the network with minimal effort.

Section 4.5 discusses the systems integration process in detail.

3.2 Proposed Network Architecture

As the design cycles progressed, we arrived at a vision for the overall structure of the network.

Hierarchical in nature, the network passes raw data up through several layers which interpret the raw data. The sensors represent the networks data acquisition capabilities. The first input into the network comes from the spectrum sensors. They scan an RF bandwidth and determine if a signal is present. If a signal is determined to be present, it is classified and passed to the geolocation sensors.

The geolocation sensors now have a set of physical layer parameters to geographically

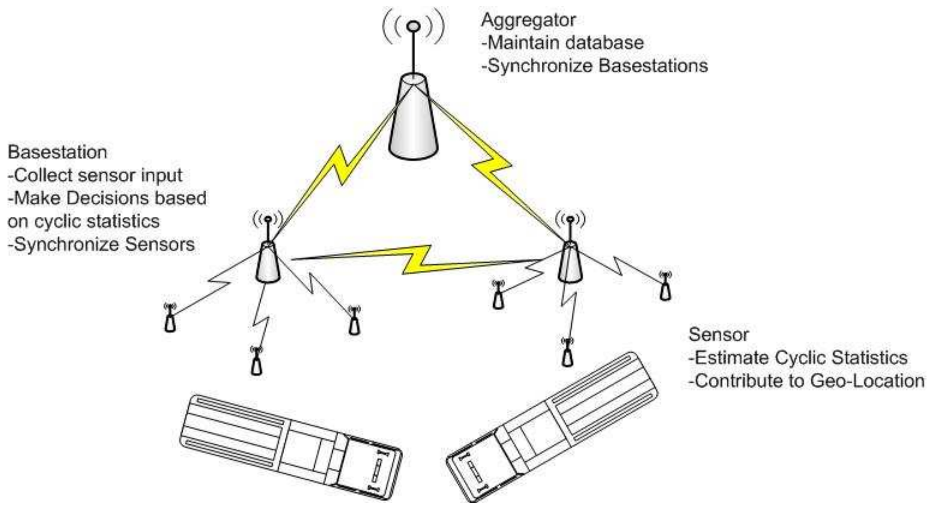

Figure 3.6: The hierarchical network architecture shown here allows for network size to be scaled up or down depending upon the size of the event. As more sensor clusters are needed, more base stations are added to handle them.

locate. The first responder communications equipment is VITA 49 compliant[1] and thus makes use of contextual packets.[2] The geolocation sensor strips out the unique ID and the time of transmission from the context packet and uses this information as well as RSSI measurements to obtain a reasonably accurate estimate of where the radio physically is.

The geolocation then the following information up the network:

1. PHY layer parameters
2. Geographical location
3. Unique ID

This information is then incorporated by the aggregator into the networks database. The base stations depicted in figure 3.6 coordinate the sensors so that they scan the same RF bandwidth and act as data fusion centers for the spectrum sensors. The sensing and classification sub-system finds and locates signals of interest. It then passes along the spectral location of these signals and instructions on how to talk to it to the rest of the network. From these parameters, the network exploits the contextual packets of the VRT standard to obtain more information about the signal of interest. The geolocation sub-system relies heavily on the sensing and classification sub-system. Since RSSI measurements are sensitive to the environment, a better way is needed to triangulate the position of a transmitter.

[1]VITA 49 is a radio transport standard for software defined radios. More information may be found at http://digitalif.org

[2]These packets are part of the VITA 49 standard and contain fields for information such as time of transmission, symbol rate, etc.

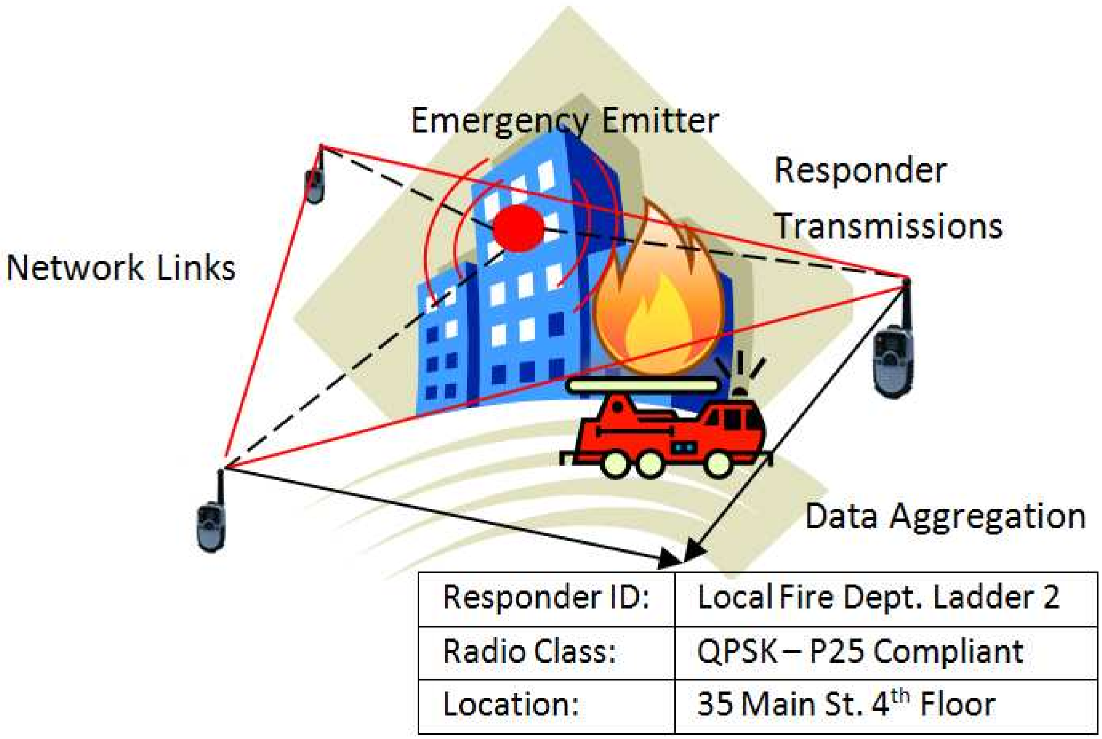

Figure 3.7: The network will locate and classify the transmissions of first responder teams.

Using the knowledge the signal's PHY parameters, it strips out the time of transmission from the contextual VRT packet. It then compares this time with the time of reception and coordinates with two other sensors to triangulate an exact position of the transmitter. The data aggregation sub-system represents the final deliverable of the network. It resides at the highest level of the network and coordinates the activities of all of the basestations. It is in charge of ensuring that each ID in the network is unique and that all responders are being persistently tracked.

3.3 Chapter Summary

An iterative design approach was taken when designing this network. This approach was chosen because of its flexibility when defining and implementing required functionality. The design process lasted several months and was ongoing until late 2009. Throughout the design, sub-systems were being designed and implemented. As individual sub-systems came together, the big picture of the network became easier to see and define. Important lessons

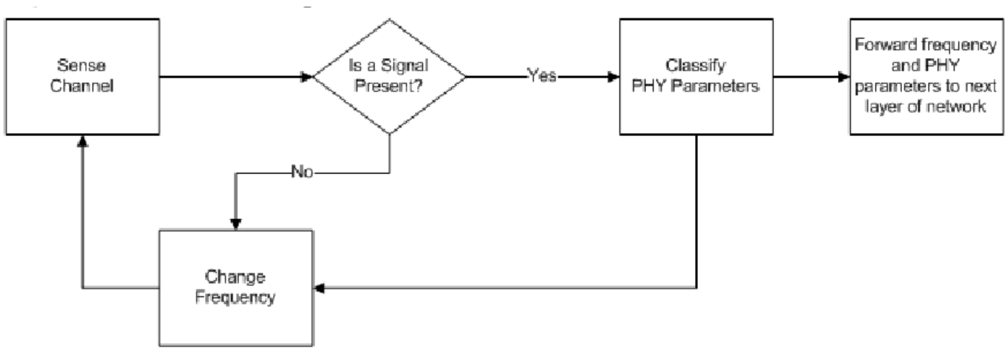

Figure 3.8: Block Diagram of the Spectrum Sensing & Signal Classification sub-systems.

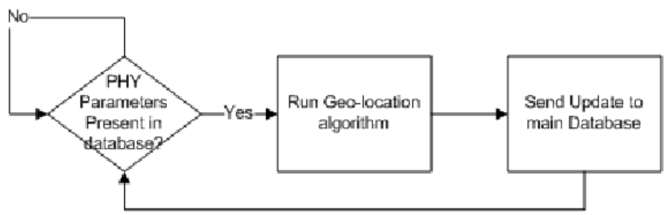

Figure 3.9: Block Diagram of the Geolocation sub-system.

learned including meticulously defining the problem statement before work begins and that working through complex mathematics is best achieved by starting from first principles.

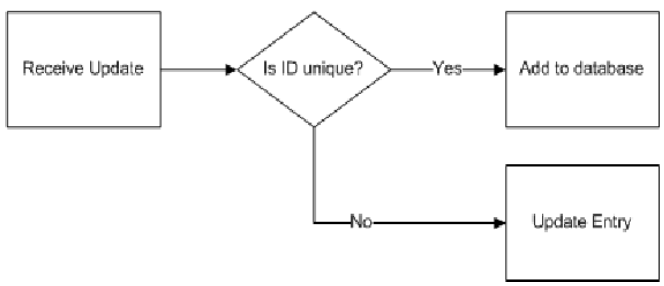

Figure 3.10: Block Diagram of the Data Aggregation sub-system.

Chapter 4

Experimental Results

4.1 A Multi-Cycle Signal Detector

The following figures illustrate some key concepts of cyclic detectors. All test vectors were generated with a simulink model similar to that in Figure 4.1.

Figures 4.2 and 4.3 served as sanity checks. Since noise is not a cyclostationary process, its expected SOF would not contain any distinct features. This is reflected in figure 4.2. The two diagonals of the figure are the autocorrelations of the noise, thus their correlation coefficients have an amplitude of 1.

Figure 4.3 shows the time domain version (SCF) of the cyclic autocorrelation function of a sinusoid. As expected, at baseband and at $\alpha = 0$ is the autocorrelation of a sin function. The other features present are at the frequency of the sinusoid (10Hz).

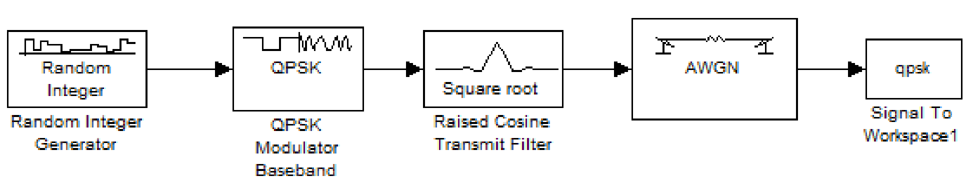

Figure 4.1: Simulink model used to generate test vectors.

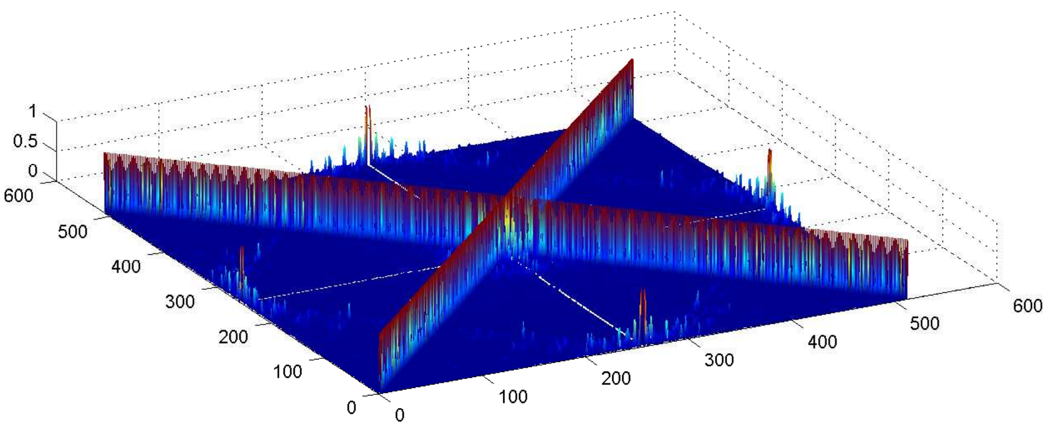

Figure 4.2: SOF of AWGN

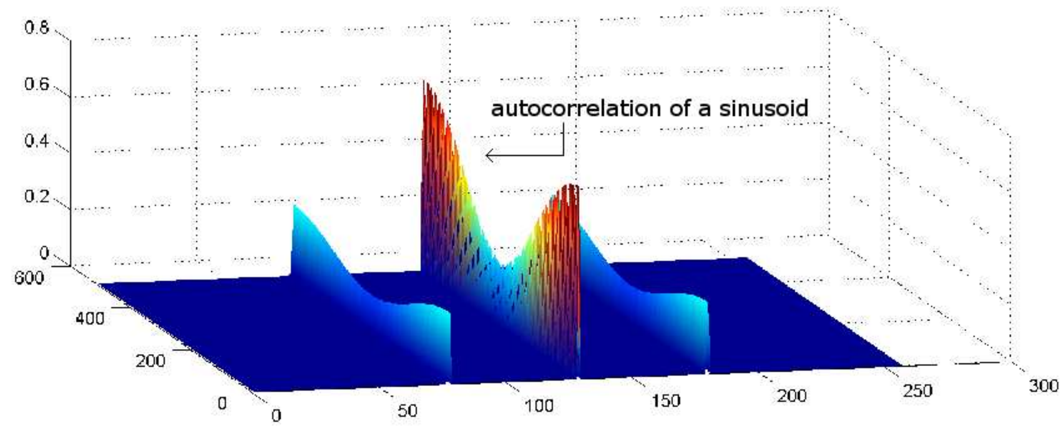

Figure 4.3: SCF of ideal 100Hz Sinusoid

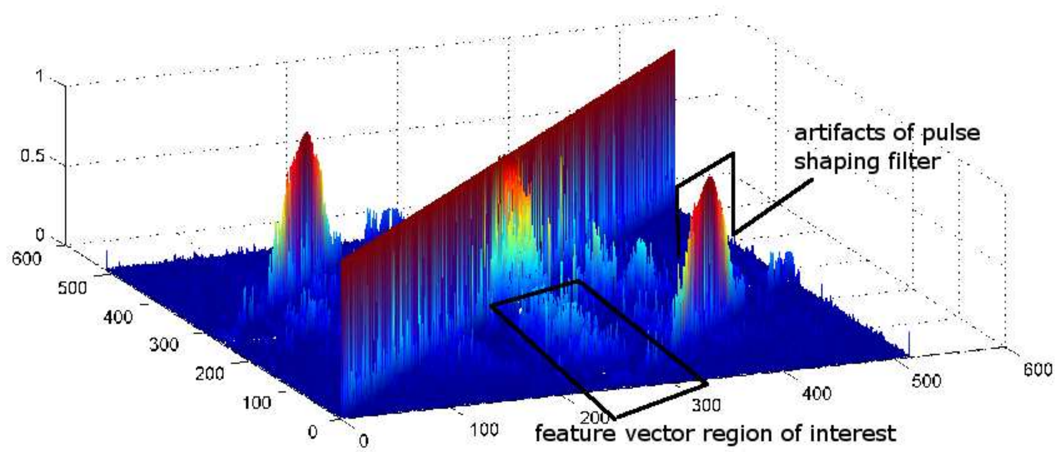

Figure 4.4: SOF of QPSK $\beta = 0.5$

Figure 4.4 shows the spectral features of an "ideal" qpsk signal. Ideal in this case means that no noise is overlapped on the transmission and the rolloff from the pulse shaping filter is high (0.5 in this case). Under these conditions, spectral features of interest should be clearly visible. Comparing Figure 4.5 to 4.4, one can observe the dampened features. This is most obvious on the pulse shaping artifact, but can still be noticed on the spectral features. At β values less than 0.3, if the channel was especially noisy, performance would decrease dramatically.

When the excess bandwidth on the pulse shaping filter is set to 0.0, all of the spectral features of interest disappear. Practically, a value of 0 would never be used, but when compared to $\beta = 0.5$, this illustrates the reliance of the cyclic detector on redundancy in the signal.

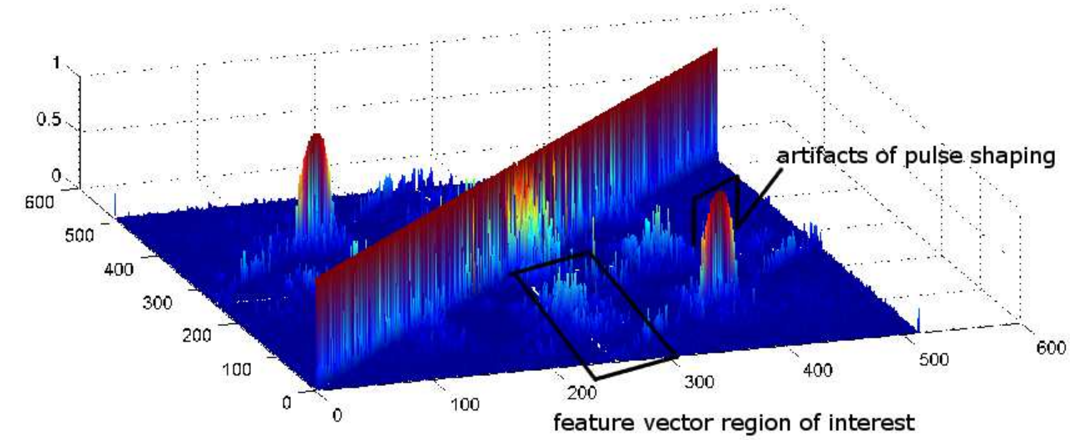

Figure 4.5: SOF of QPSK $\beta = 0.3$

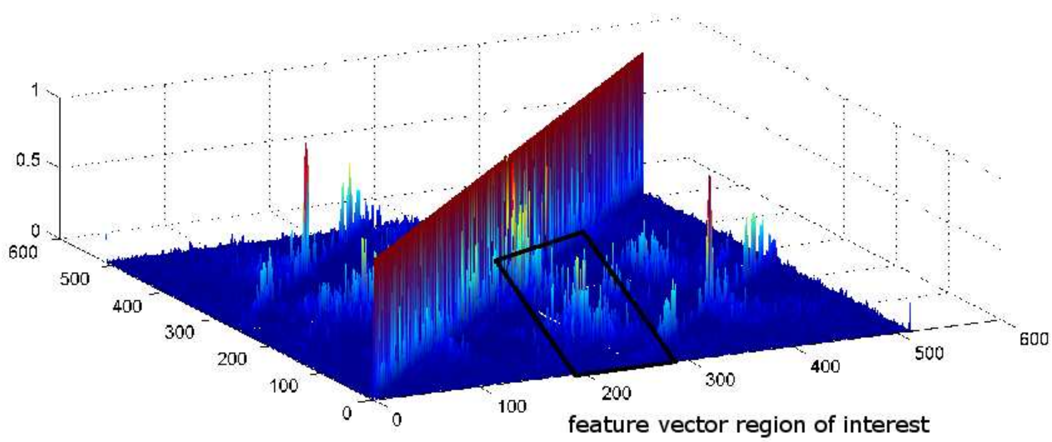

Figure 4.6: SOF of QPSK $\beta = 0$

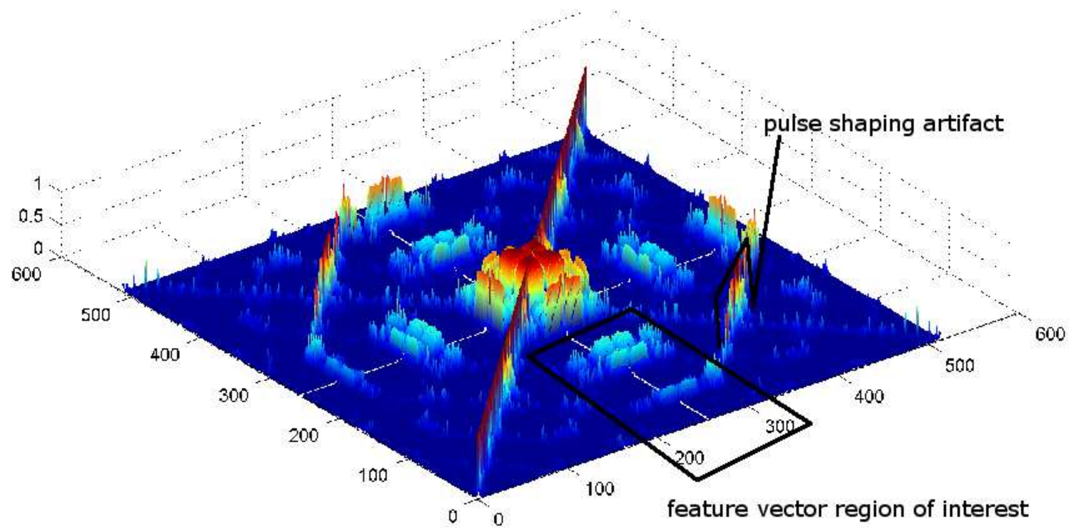

Figure 4.7: SOF of 4-FSK $\beta = 0.3$

Figures 4.4, 4.5, and 4.6 show the SOF of QPSK signals with varying pulse shapes. Though a standard square root raised cosine filter is used on the transmitter for each vector, the rolloff factor, β was varied at each step. This also verifies the intuitive conclusion of cyclic analysis: the more redundancy in your data, the more obvious cyclic features will be. In this case, the redundancy takes the form of the excess bandwidth typically used to guard against intersymbol interference.

Figures 4.8 and 4.7 show the SOF of 4-FSK signals, again with varying pulse shapes. It is clear that the SOFs of FSK and QPSK are quite distinct. This uniqueness makes signal classification seem an almost automatic next step in the analysis of these signals.

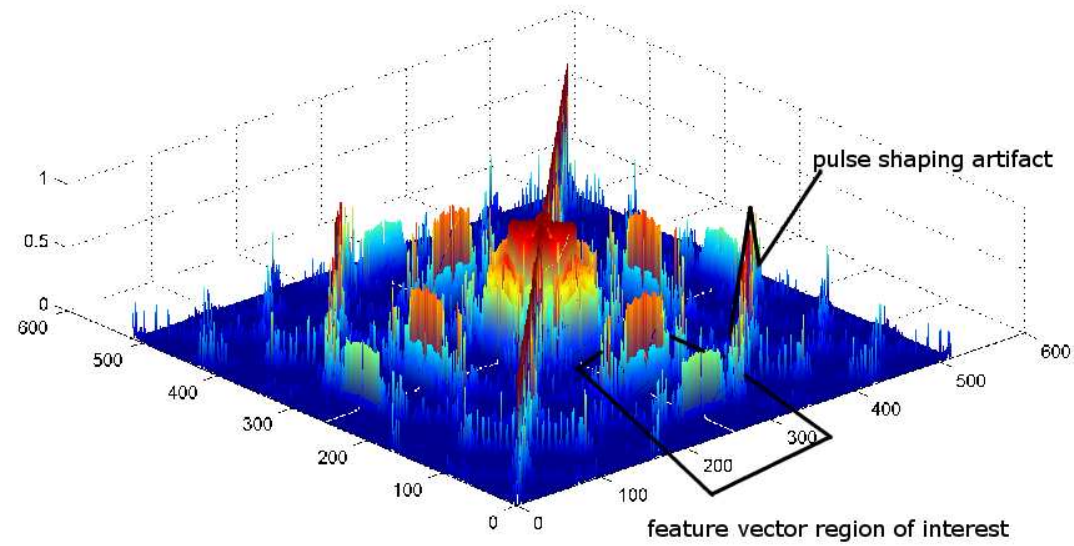

Figure 4.8: SOF of 4-FSK $\beta = 0.5$

4.2 Signal Classification

Signal classification was accomplished by taking a profile of the alpha domain given in [16] by:

$$I(\alpha) = \max_f |C_x^\alpha(f)| \tag{4.1}$$

The key take away from Figures 4.9 and 4.10 are that they are so different. There is no way that one could be mistaken for the other. These profiles must be evaluated during the network initialization period, as they are unique to the channel.

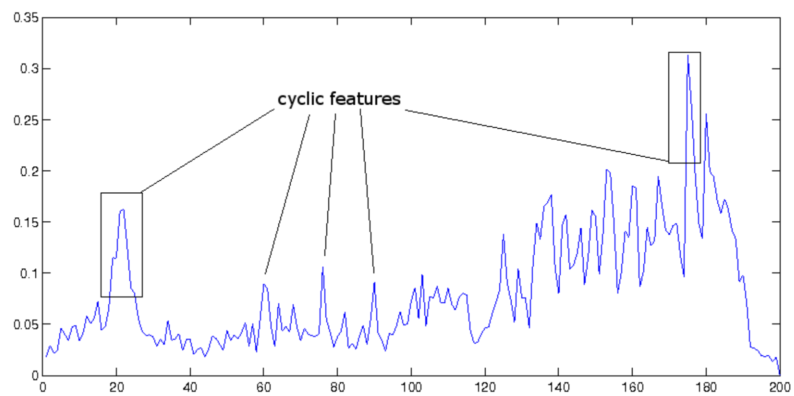

Figure 4.9: α profile of QPSK

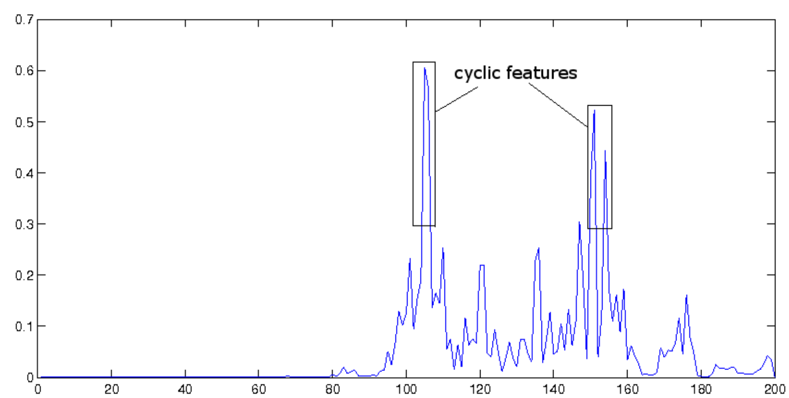

Figure 4.10: α profile of 4-FSK

Figures 4.10 and 4.9 show the profile calculated for ideal QPSK and 4-FSK signals. Ideal here means that the signals were uncorrupted by noise and were transmitted through an SRRC filter with $\beta = 0.5$. Their profiles were calculated from an SOF estimated from 1000

Table 4.1: Signal Classification - QPSK at 15dB SNR

Trial	QPSK Classified	FSK Classified	Noise Classified
1	6	0	4
2	7	0	3
3	6	0	4
4	5	0	5
5	6	1	2

Table 4.2: Signal Classification - 4-FSK at 15dB SNR

Trial	QPSK Classified	FSK Classified	Noise Classified
1	2	6	2
2	0	7	3
3	2	5	3
4	0	7	3
5	1	7	2

samples with a windowing size of 512 samples. The profiles used are also the average of 5 "ideal" signals.

These profiles exist within the network before it is initialized.

[16] refers to this profile as the Cyclic Domain Profile (CDP). To determine if a signal is present, the correlation coefficient described in Equation 2.11 between the ideal CDPs and observed data samples is calculated. The CDP for 4-FSK and QPSK are different enough such that a strong correlation against the ideal 4-FSK CDP will not correlate strongly to the QPSK CDP. This means if you compute the correlation coefficient of an FSK signal vector with the ideal CDP of QPSK, the value is very low (on the order of 0.1-0.3). However computing the coefficient between the FSK vector and the ideal FSK CDP typically yields values upwards of 0.6-0.7.

Once a comparison made and the closest match selected, the coefficient is compared against a threshold to differentiate a data signal from noise. This threshold is computed by evaluating ρ_{FSK} and ρ_{QPSK} for the channel noise. These thresholds can also be reevaluated in environments where channel conditions might vary with time.

4.3 Data Fusion

The preferred data fusion solution chosen for this network is described by Section 2.5.2. Table 4.3 shows several sample data sets and how the confidence has been improved.

Table 4.3: Data Fusion Algorithm Results

Input Vector	Reliability Vector	Final Decision	Confidence
[0 0 2 2 2 2 2 1 1 0]	[0.2 0.3 0.7 0.6 0.7 0.7 0.7 0.4 0.3 0.1]	[2]	[0.71]
[0 0 2 2 2 2 2 1 1 0]	[0.8 0.7 0.2 0.3 0.4 0.2 0.1 0.4 0.4 0.9]	[0]	[0.54]
[0 0 2 2 2 2 2 1 1 0]	[0.8 0.7 0.2 0.3 0.9 0.9 0.1 0.4 0.4 0.9]	[2]	[0.42]
[0 0 2 2 2 2 2 1 1 0]	[0.2 0.3 0.7 0.6 0.7 0.7 0.7 0.4 0.9 0.1]	[2]	[0.64]

4.4 Verification in Hardware

Though success was demonstrated in software, a hardware metric was sought to validate theoretical performance. The WPI Wireless Innovation Lab has developed a prototype interface between Simulink & the Universal Software Radio Peripheral (USRP) 2.0 SDR platform [19].

Test vectors were created in Simulink and transmitted over the air. The received vector was then put through the same algorithm as the purely theoretical vector. Figure 4.11 show the set up used to collect data.

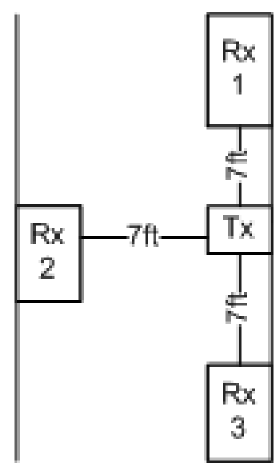

Figure 4.11: Configuration of hardware verification experiment. This experiment was conducted on the second floor of Atwater Kent at WPI. Four USRP 2.0s were used: 3 as receivers, 1 as a transmitter.

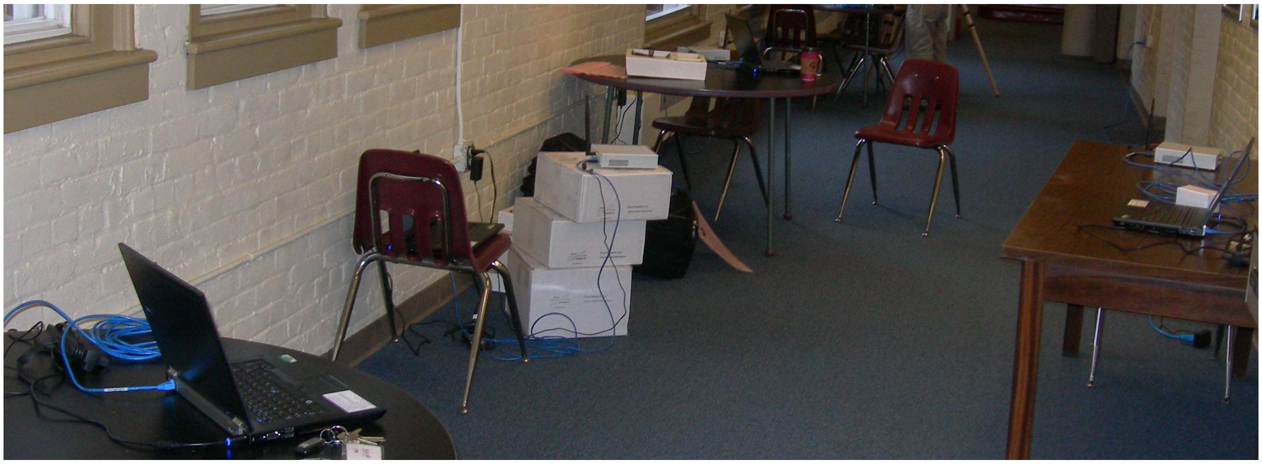

Figure 4.12: Photograph of the actual experimental setup used to verify detection algorithms and SDR -¿ Simulink interface.

Figure 4.13: Two USRP 2.0 radios set up to transmit and receive in the WPI Wireless Innovation Lab.

The resulting SOF (given by Equation (2.11)) looked like figure 4.14. Non of the expected spectral features are present, however this does not mean there is a problem with the algorithm. The interface used to transmit and receive this data is severely constrained by the local computers it is run on. Likely, the laptops were too slow to process all of the data sequentially and packets were dropped.

This test served a dual purpose of testing the classification algorithm (which correctly recognized the transmission as noise), and benchmarking the SDR interface under development by the WPI Wireless Innovation Laboratory.

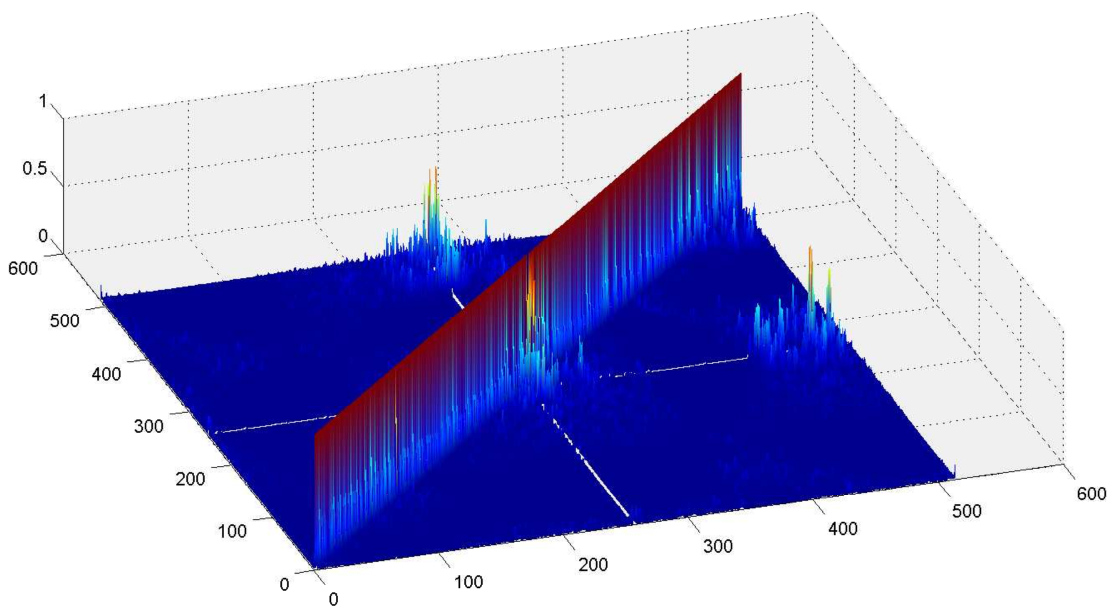

Figure 4.14: SOF of QPSK as Transmitted OTA

4.5 Systems Integration

Since these sub-systems will be functional blocks in a larger operational network, user-friendliness was considered. The functionality described here was implemented as a black box. Figure 4.15 shows exactly where the sub-systems implemented in this report fall into the entire network.

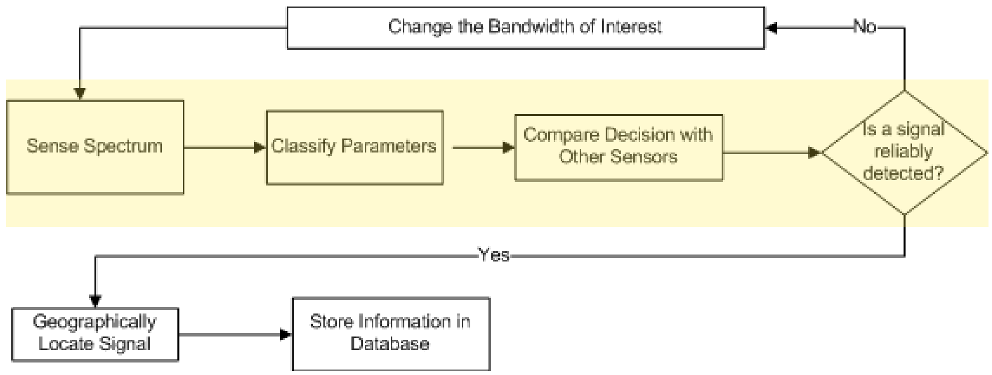

Figure 4.15: A block diagram of the data flow throughout the entire network.

Function	Inputs	Outputs
Sensing & Classification	1000 data samples	signal type
	FSK ideal CDP	
	QPSK ideal CDP	
	FSK Detection Threshold	
	QPSK Detection Threshold	
Data Fusion	Vector of N decisions	final decision
	Vector of N reliability estimates	confidence

Functions that required an input listed their input and output parameters. Such parameterization combined with the known purpose of the block provides an effective black box set up that ensures functionality without demanding tweaking or maintenance.

In the event additional input types need to be considered, the code is well commented. Appendix A and B of this report contain the source code for the classification/sensing & data fusion blocks respectively.

4.6 Chapter Summary

Both signal detection & classification and data fusion sub-systems were successfully implemented. The cyclic detector can successfully detect and classify a signal with a 60-70% success rate in a low SNR environment. The data fusion module is capable of taking several sensor inputs and interpreting a single, reliable decision from all of them. The data fusion module improves the detection capabilities of the network significantly.

Chapter 5

Conclusions

Cooperative signal detection and classification are important sub-systems in modern communications networks. With direct applications to dynamic spectrum access, decentralized networks, synchronization, and signals intelligence, such sub-systems can increase network capacity, reduce bit error rates, and optimize spectrum usage.

The implementation described in this report has successfully achieved the following goals:

1. Spectrum Sensing: Signals were successfully detected in unfavorable channels. No prior knowledge was assumed about these signals.
2. Signal Classification: Once detected, signals were reliably classified with over a 60% success rate.
3. Data Fusion: Decisions of several sensors were combined to increase the reliability of successful detection to over 80%.

The network that these sub-systems will be utilized in represents an important step toward realizing the next generation of first responder communications. The network will demonstrate the potential of SDR and cognitive radio technology to first responders. One of the desired outcomes of the SRC 2009 is a network that could be translated into an actual implementation.

By demonstrating the practical feasibility of these networks, the foundation is being established for a true implementation to be proposed. Standards like VRT and Project 25 are standarizing the hardware and infrastructure that first responders rely upon. With common equipment, it is only a matter of time before a network like the one described in this report is implemented.

Interoperability is a recognized need in the first responder community. With industry facilitating the implementation of standards to encourage this goal, the sub-systems implemented in this report represent an important first step to demonstrating the potential and proving the reliability of networks that take advantage of SDR technology.

5.1 Future Work

Despite these successes, there are many intuitive extensions of this work that could provide challenging undergraduate or graduate level research projects.

- The cyclostationary signal detector described in this report was tested against AWGN channels. Though success was demonstrated in low SNR channels, fading channels present a unique challenge to cyclic detectors. Fading can directly suppress spectral features in signals. Implementing adaptive channel equalization prior to the cyclic detector might optimize the cyclic detectors operating in fading channels. This would be a challenging problem, as additional computational trade offs may be needed to achieve real-time operation.
- Develop an efficient search algorithm for a real-time cycle detector. Such an algorithm would need to traverse a large RF bandwidth of interest searching for indications of modulated data with no prior knowledge of the transmissions, center frequencies, or frequency hopping pattern.
- The well-defined scope of the project simplified the signal classifier. Since QPSK and FSK have distinct spectral features, a basic comparison was sufficient to differentiate the modulation schemes. In more complex applications, a signal detector implemented as a trained network would be more robust to false detection and have expanded functionality.
- Since the channels were primarily AWGN in nature, and there were few possible decisions, the data fusion center was also simplistic in nature. If a signal classifier were trained to look for a large range of signals, Bayesian filtering would offer a more optimal solution than the weighted averaging method described in this report.

  Though there is a lot of opportunity for expansion on this implementation, the sub-systems described here were sufficient to demonstrate several important points:

- Computational power has achieved levels required to implement real-time cyclic detectors with a reasonable frequency resolution. Though [6] described such systems in detail almost two decades ago, only recently have implementations and applications become feasible.

- Cyclic detectors can detect a signal in harsh channel environments - especially if they want to be found. Such success is important because systems described in [15] that rely on envelope analysis have probabilities of false detection.
- signal classification is a natural and intuitive extension of cyclic detectors.

Such sub-systems will play important roles in modern wireless communications. As consumer, military, and industrial applications are developed for point-to-point ad hoc networks, they will rely upon such sub-systems to optimize limited spectrum usage, reduce overhead, and encourage decentralized networks.

Bibliography

[1] T. H. Kean and L. H. Hamilton, "9/11 comission report," 2004.

[2] P. G. Cook and W. Bonser, "Architectual overview of the speakeasy system," *IEEE Journal on Selected Areas in Communications*, 1999.

[3] C. R. C. M. da Silva, B. Choi, and K. Kim, "Distributed spectrum sensing for cognitive radio systems," Virginia Polytechnic Institute and State University, Tech. Rep., 2007.

[4] "Smart radio challenge 2009," Software Defined Radio Forum, 2009, http://radiochallenge.org/09SampleProblem.html.

[5] D. Hall, *Mathematical techniques in multi-sensor data fusion.* Artech Print, 2004.

[6] W. A. Gardner, "Exploitation of spectral redundancy in cyclostationary signals," *IEEE Signal Processing Magazine*, 1991.

[7] E. Like, "Non-cooperative modulation recognition via exploitation of cyclic statistics," Master's thesis, Wright State University, 2007.

[8] K. Po and J. Takada, "Signal detection method based on cyclostationarity for cognitive radio," Tokyo Institute of Technology, Tech. Rep., 2007.

[9] *The VITA Radio Transport as a Framework for Software Definable Radio Architectures.* SDR 08 Technical Conference and Product Exposition, 2008.

[10] O. Filth, image created by Oli Filth, using Matlab and Adobe Photoshop.

[11] "Project 25," Association of Public Safety Communications Officials International, http://www.apco911.org/frequency/project25/information.html.

[12] M. Rice, *Digital Communications: A Discrete-Time Approach.* Upper Saddle River: Pearson/Prentice Hall, 2009.

[13] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 25, 2005.

[14] "802.22," IEEE 802 LAN/MAN Standards Committee 802.22 WG on WRANs, http://www.ieee802.org/22.

[15] D. Datla, "Spectrum surveying for dynamic spectrum access networks," Master's thesis, University of Kansas, 2006.

[16] K. Kim, I. A. Akbar, K. K. Bae, J. Um, C. M. Spooner, and J. H. Reed, "Cyclostationary approaches to signal detection and classification in cognitive radio," Virginia Polytechnic Institute and State University, Tech. Rep., 2007.

[17] G. Welch and G. Bishop, "An introduction to the kalman filter," University of North Carolina at Chapel Hill, Tech. Rep., 2006.

[18] R. S. Roberts, W. A. Brown, and J. Herschel H. Loomis, "Computationally efficient algorithms for cyclic spectral analysis," *IEEE Signal Processing Magazine*, 1991.

[19] M. J. Leferman and A. M. Wyglinski, "Taming software-defined radio: A graphical user interface for digital communication system prototyping," *DSP-FPGA.com*, 2010.

Appendix A

Code: Signal Detector & Classifier

```
%#eml
function [S Cx f_profile signal con] = get_SCF(x,fski,qpski,Cthfsk,Cthqpsk)
% get_SCF
%  The HOSA toolbox & Cyclostationary Toolboxes provided a reference point
%  for some of this code.  Both are no longer supported.
%              Input parameters:
%              x      signal
%              fski   ideal vector for FSKi - initalize as something random
%              and then pre-compute
%              qpski  ideal vector for QPSK profile - intialize as
%              something random, and then pre-compute
%              Cthfsk - Threshold for FSK detection
%              Cthqpsk - Threshold for QPSK detection
%
x = x';
y = x;
%N must be even and divisible by 4 and < lx
N = 512;
lx=length(x);
%If no Cth specified, default to these
%These should be tweaked during network start up
%for best performance.  The worse the channel, the lower they should be.
%In an ideal channel, they should be closer to 0.8
if(nargin==3)
    Cthfsk = 0.5;
    Cthqpsk = 0.5;
end

%Set up variables
n=0:floor(lx-N);
ln=length(n);
```

```
%Compute windowing functions for later.
a=feval('hamming',N)';
g=feval('hamming',ln)';
g=g/sum(g);
a=a/sum(a);
Ts=1/N;
%Pre-allocate for speed
S=zeros(N+1,N/2+1);
X=zeros(2*N+1,ln);
Y=zeros(2*N+1,ln);

%Freq. Smoothed Cyclic Periodogram
for f=-N:N
  %N point FFTs of the signal are computed
  xf=x.*exp(-1i*2*pi*f*(0:lx-1)*Ts);
  yf=y.*exp(-1i*2*pi*f*(0:lx-1)*Ts);
  for i=1:ln
    %Multiply the FFT of X with the conj of Y and vice versa
    n_r=n(i)+(1:N);
    X(f+N+1,i)=a*xf(n_r)';
    Y(f+N+1,i)=conj(a*yf(n_r)');
  end
end

for alpha=-N/4:N/4
  for f=-N/2:N/2
    f1=f+alpha;
    f2=f-alpha;
    if ((abs(f1)<N/2)&&(abs(f2)<N/2))
      %g acts to smooth X*Y out, this is more obvious if you plot g
      %s is the cross correlation of X's and Y's frequency components
      %seperated by f +/- alpha
      S(f+N/2+1,N/4+alpha+1)=g*(X(f1+N+1,:).*Y(f2+N+1,:))';
    end
  end
end
%External EML function
eml.extrinsic('corrcoef');

%Compute correlation coefficients
Cx = fftshift(corrcoef(S').^2);
%Extract feature vector region of interest
features = abs(Cx(1:200,240:280));
f_profile = zeros(1,200);
%Take a snapshot of the most outstanding features in the region.
```

```
for f = 1:length(features(:,1))-1
    f_profile(f)=max(features(f,:));
end
%Start Classification
%Determine which profile most closely matches the signal.
a = abs(corrcoef(f_profile,fski));
b = abs(corrcoef(f_profile,qpski));
%Now determine if that match is sufficient to deem the signal present.
if a(1,2) >= b(1,2)
    st = 1;
else
    st = 2;
end
switch st
    case 1
        if a(1,2) >= Cthfsk
            signal = 1;
            con = a(1,2);
        else
            %If the threshold is not met, determine the signal absent
            signal = 0;
            con = a(1,2);
        end
    case 2
        if b(1,2) >= Cthqpsk
            signal = 2;
            con = b(1,2);
        else
            %If the threshold is not met, determine the signal absent
            signal = 0;
            con = b(1,2);
        end
end
end
```

Appendix B

Code: Data Fusion

```
%Data Fusion
function [final con] = fusion(input, rel)

linput = length(input);
lrel = length(rel);

%input is a vector of length linput with potenial values:
% 0 - noise
% 1 - 4-FSK
% 2 - QPSK

if(lrel ~= linput)
    disp('rel and input must be same length');
    return;
end
%Indexs of decisions
votenoise = find(input == 0);
votefsk = find(input == 1);
voteqpsk = find(input == 2);

%Sum confidence of sensors
noise = sum(rel(votenoise));
fsk = sum(rel(votefsk));
qpsk = sum(rel(voteqpsk));

%Compute total amount of confidence in the system
tot = noise+fsk+qpsk;

%Weight the number of votes by respective confidence
n = length(votenoise)*noise;
f = length(votefsk)*fsk;
```

```
q = length(voteqpsk)*qpsk;

%Determine output
%confidence is determined by determinig the ratio of
%confidence in 1 decision to the total confidence
%con represents the confidence in the final decision.
% 0 1 or 2 represent noise, 4-fsk, and qpsk respectively
[~,I] = max([n f q]);
switch I
    case 1
        final = 0;
        con = noise / tot;
    case 2
        final = 1;
        con = fsk / tot;
    case 3
        final = 2;
        con = qpsk/tot;
end
end
```